U.S. Coast Guard Research and Development Center
1082 Shennecossett Road, Groton, CT 06340-6096

Report No. CG-D-15-01

REVIEW OF SEARCH THEORY: ADVANCES AND APPLICATIONS TO SEARCH AND RESCUE DECISION SUPPORT

FINAL REPORT
SEPTEMBER 2001

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

This document is available to the U.S. public through the National Technical Information Service, Springfield, VA 22161

Prepared for:

U.S. Department of Transportation
United States Coast Guard
Operations (G-OPR)
Washington, DC 20593-0001

Technical Report Documentation Page

1. Report No.	2. Government Accession Number	3. Recipient's Catalog No.
CG-D-15-01		

4. Title and Subtitle Review of Search Theory: Advances and Applications to Search and Rescue Decision Support	5. Report Date September 2001 6. Performing Organization Code Project No. 1012.3.3
7. Author(s) J. R. Frost (Soza & Co.) L. D. Stone (Metron, Inc.)	8. Performing Organization Report No. R&DC 554
9. Performing Organization Name and Address Soza & Company, Ltd. 8550 Arlington Blvd. Fairfax, VA 22031 Metron, Inc. 11911Freedom Dr. Suite 800 Reston, VA 20190 U.S. Coast Guard Research & Development Center 1082 Shennecossett Road Groton, CT 06340-6096	10. Work Unit No. (TRAIS) 11. Contract or Grant No.
12. Sponsoring Organization Name and Address U.S. Department of Transportation United States Coast Guard Office of Operations Washington, DC 20593-0001	13. Type of Report & Period Covered Final 14. Sponsoring Agency Code Commandant (G-OPR) U.S. Coast Guard Headquarters Washington, DC 20593-0001

15. Supplementary Notes

The R&D Center's technical point of contact is Dr. Jennifer Dick O'Donnell, 860-441-2741, email: jodonnell@rdc.uscg.mil.

16. Abstract (MAXIMUM 200 WORDS)

Fundamental limitations inherent in manual search planning methods have severely limited the application of advances in several areas that could improve the efficiency and effectiveness of the U.S. Coast Guard's search and rescue mission. These areas include advances in search theory, environmental data products, knowledge of detection profiles for various sensors, and knowledge of leeway behavior. The U.S. Coast Guard's computerized search planning aids have not kept up with advances in these areas or with technology in general. This report reviews the history and recent advances of search theory and its application to a variety of search problems. It then reviews the history of the U.S. Coast Guard's search planning methods, showing where search theory was initially applied, albeit in a necessarily very limited way, and where later modifications departed from the theoretical basis of the original methodology. Several computerized search planning decision support tools are analyzed and compared, as are the differences between an analytic approach and a simulation approach. The results are summarized in a matrix.

The U.S. Coast Guard needs a new search planning decision support tool for search and rescue and other missions. This tool should use the simulation approach due to its power and flexibility as compared to analytic techniques.

17. Key Words		18. Distribution Statement
search and rescue	search planning	This document is available to the U.S. public through the National Technical Information Service, Springfield, VA 22161
search theory	search objects	
SAR planning	SAR simulation	

19. Security Class (This Report)	20. Security Class (This Page)	21. No of Pages	22. Price
UNCLASSIFIED	UNCLASSIFIED	127	

Form DOT F 1700.7 (8/72) Reproduction of form and completed page is authorized.

EXECUTIVE SUMMARY

This report describes the developments in the field of search theory from its origins during World War II to the present day. The completion of this report coincides with the formation of a team to develop the Coast Guard's (CG) new search planning tool, thus providing a perfect opportunity to use this knowledge to develop a search tool that can handle more complex search scenarios with more accurate predictive capabilities. The compromises, simplifications and inaccuracies that have been introduced into the CG search theory over time are described in this report. This report also presents recommendations for the form and functionality of the CG's new search planning tool. These findings can provide the basis for the CG's new search theory tool.

The basic principles of search theory were first applied to SAR planning around 1957 when the U.S. Coast Guard published its search planning doctrine in a search and rescue manual. To apply search theory to practical SAR planning problems, since computers were not then widely available, simplifications to the theory had to be made to develop a method feasible for hand calculations. This became known as the "classical search planning method" (CSPM), and it remains the basis for search planning support tools today.

CSPM was originally a scientifically based, analytic method that was appropriate for the technology and data available when it was implemented in 1957. However, due to the technological limitations, it could address only the simplest of SAR scenarios. Later, attempts were made to extend the methodology to more typical, and more complex, situations. Sometime after 1963, a number of sub-optimal "field modifications" were made that are inconsistent with the underlying theory. The most notorious of these modifications is the Min/Max technique, which violates some of the basic assumptions of the CSPM and the scientific principles on which it was based. Partial attempts to rectify this situation, like the mid-point compromise, were not always improvements.

The Search and Rescue Planning (SARP) system, the first implementation of search planning support on computers, occurred around 1970, well before the microcomputer age. SARP was basically a computerized version of the then current version of the modified CSPM with somewhat improved use of environmental data and drift computations. A few years later SARP was joined by the Computer Assisted Search Planning (CASP) system that took a computer simulation approach to the search planning and evaluation problem. Unfortunately, Coast Guard search planning support tools have not kept pace with technological advances in three important respects. First, they have not kept up with advances in search theory and algorithm development relevant to the practical application of search theory using computer simulation. Secondly, modifications made to the CSPM to make it applicable to typical, complex search scenarios are inconsistent with the basic theory. Third, and most importantly, the tools have generally not kept pace with the significant increases in the amount, level of detail, or accuracy of environmental data products or with new knowledge about drift behavior or detection.

These shortcomings are evident in the four current computerized SAR planning support tools. Their primary difference from SARP is that most run on desktop computers and are supported

by geographic information systems and other interactive tools. However, they all suffer from the limitations of the CSPM. Ironically, the Coast Guard's C2PC/SAR Tools Automated Manual Solution suffers the most in this regard because it is the most faithful to the manual technique. The other three similar tools provide considerable improvement in the availability and use of environmental data for drift computations. CASP 1.x is the most capable search planning tool available today, but only because it is the one tool that uses simulation techniques. CASP's basic framework is more than 25 years old and even then it was severely limited by the obsolete computing environment in which it was forced to operate. In short, CASP 1.x is a primitive implementation of a sophisticated methodology and many key elements are still missing.

In more recent years, the greatly increased availability of inexpensive but powerful microcomputers with high-resolution color geographic information systems has made near-real-time computing support a normal expectation. Programming powerful workstations to emulate the hand calculations of the manual method, a technique developed to get around the lack of computing capability in the 1950's, is unacceptable. Presently, the Coast Guard currently uses more sophisticated techniques to predict oil spill trajectories and perform risk analyses than for matters of life and death, such as SAR.

It is recommended that the Coast Guard's SAR planning theory be corrected and that stochastic analysis be the primary method for executing search plans. The U. S. Coast Guard needs and deserves a new computer simulation-based search planning support tool that takes full advantage of the advances in these areas to ensure efficient, effective use of expensive search resources. Those awaiting rescue deserve the time advantage such a planning support tool can offer.

TABLE OF CONTENTS

Page

EXECUTIVE SUMMARY ... iv

CHAPTER 1. INTRODUCTION ... 1-1

1.1 BACKGROUND ... 1-1

1.2 OVERVIEW ... 1-2

CHAPTER 2. THE DEVELOPMENT OF SEARCH THEORY ... 2-1

2.1 INITIAL DEVELOPMENT ... 2-1

2.2 TYPES OF SEARCH PROBLEMS ... 2-2

2.2.1 One-Sided Search Problems ... 2-2

2.2.1.1 Stationary Search Objects ... 2-3

2.2.1.2 Moving Search Objects ... 2-3

2.2.2 Two-sided Search Problems ... 2-3

2.2.2.1 Cooperative Searches ... 2-3

2.2.2.2 Non-Cooperative Searches ... 2-4

2.3 KOOPMAN'S EARLY WORKS ... 2-4

2.3.1 Effective Sweep Width ... 2-4

2.3.2 Exponential Detection Function ... 2-5

2.3.3 Search Object Location Distribution ... 2-6

2.3.4 Optimal Search Effort Density ... 2-7

2.3.5 Posterior Search Object Location Density ... 2-8

2.4 NON-EXPONENTIAL DETECTION FUNCTIONS AND CELLULAR DISTRIBUTIONS ... 2-8

2.4.1 Non-exponential Detection Functions ... 2-9

2.4.2 Optimal Allocations for Cellular Distributions ... 2-10

2.4.3 Uniformly Optimal Search Plans ... 2-11

2.5 OPTIMAL SEARCH FOR MOVING AND MULTI-STATE SEARCH OBJECTS ... 2-11

2.5.1 Optimal Search for Moving Search Objects ... 2-11

2.5.2 Optimal Survivor Search 2-12

2.6 CONSTRAINTS ON THE SEARCHER 2-12

2.6.1 Path Constraints 2-12

2.6.2 Simplicity Constraints 2-12

2.7 SEARCH AND EVASION PROBLEMS 2-14

CHAPTER 3. PRACTICAL APPLICATIONS OF SEARCH THEORY 3-1

3.1 RELATING THEORY AND PRACTICE 3-1

3.2 NAVAL APPLICATIONS DURING WORLD WAR II 3-2

3.3 MORE RECENT NAVAL APPLICATIONS 3-4

3.3.1 Finding Lost Objects 3-4

3.3.2 Soviet Submarines 3-4

3.3.3 Clearance 3-5

3.4 HISTORICAL WRECKS—*SS CENTRAL AMERICA* 3-6

3.5 THE EFFICACY OF THEORY-BASED SEARCH PLANNING TECHNIQUES 3-9

CHAPTER 4. CLASSICAL SEARCH PLANNING METHOD 4-1

4.1 INTRODUCTION 4-1

4.2 PRIOR PROBABILITY DENSITY DISTRIBUTIONS 4-2

4.2.1 Incident Position 4-2

4.2.2 Post-incident Search Object Motion and Datum 4-2

4.2.3 Applicable Theorems from Statistics 4-3

4.2.4 Computing Datum 4-4

4.3 RELATING *POD* AND EFFORT DENSITY (COVERAGE) 4-6

4.3.1 Koopman's "Inverse Cube" Model of Visual Detection 4-6

4.3.1.1 Instantaneous Visual Detection Probability 4-7

4.3.1.2 Lateral Range Detection Function 4-8

4.3.1.3 The POD vs. Effort Density (Coverage) Detection Function 4-9

4.4 CONSTRAINTS ON EFFORT 4-11

4.4.1 Constraints on *Effort* and *Search Effort* 4-12

4.4.2 Other Constraints on *Effort* and *Search Effort* 4-12

4.4.3 Accounting for Constraints on *Effort* and *Search Effort* 4-13

4.5 OPTIMAL EFFORT ALLOCATION 4-13

4.5.1 Koopman's Uniform Coverage Squares 4-13

4.5.2 Koopman vs. Classical Search Planning "Safety Factors" 4-16

4.5.3 Optimal Search Radius for Maximum POS 4-17

4.5.4 Rationales for the CSPM First Search Radius.. 4-18

4.5.5 Subsequent CSPM "Safety Factors" 4-19

4.6 CLASSICAL SEARCH PLANNING IN PERSPECTIVE... 4-20

4.6.1 Pre-Search Probability Density Distributions 4-21

4.6.1.1 Incident Position *4-21*

4.6.1.2 Mean Drift Estimation *4-22*

4.6.1.3 Drift Error Estimation *4-24*

4.6.1.4 Summary of Pre-Search Distribution Issues.. *4-32*

4.6.2 The Detection Function 4-32

4.6.2.1 Inverse Cube Model of Instantaneous Visual Detection................................ *4-33*

4.6.2.2 Search Platform and Search Object Motions *4-34*

4.6.2.3 Summary of Inverse Cube Issues... *4-38*

4.6.3 Post-Search Probability Density Distributions.. 4-39

CHAPTER 5. MIN/MAX MODIFICATIONS TO CLASSICAL SEARCH PLANNING....... 5-1

5.1 INTRODUCTION... 5-1

5.2 THE "MINIMAX" TECHNIQUE 5-1

5.2.1 LEEWAY RATE 5-1

5.2.2 Integrating Min/Max with Classical Search Planning... 5-5

5.2.3 Leeway Divergence 5-6

5.3 QUESTIONS ABOUT THE LOGIC OF MIN/MAX.. 5-9

CHAPTER 6. COMPARISON OF SEARCH PLANNING AND CASE MANAGEMENT TOOLS 6-1

6.1 INTRODUCTION... 6-1

6.2 MANUAL SEARCH PLANNING METHODS 6-1

6.3 AUTOMATED MANUAL METHODS 6-3
6.3.1 USCG Search and Rescue Planning (SARP) 6-4
6.3.2 CANSARP 6-7
6.3.3 *ASA* SARMAP/ARCVIEW® 6-9
6.3.4 Search and Rescue Information System (SARIS) with UK CG3 Method 6-11
6.3.5 USCG C2PC Automated Manual Solution (AMS) 6-15
6.3.6 Summary of Automated Manual Solutions 6-19
6.4 STOCHASTIC COMPUTER SIMULATION 6-20
6.4.1 Monte Carlo Modeling (CASP 1.x) 6-21
6.4.2 CASP 2.0 6-26
6.4.3 Other Simulation Techniques 6-27
6.4.4 Matrix Summary of Search Planning Tools 6-27
6.5 OTHER SEARCH PLANNING DECISION SUPPORT TOOLS 6-30
6.6 CASE MANAGEMENT TOOLS 6-30
CHAPTER 7. CONCLUSIONS AND RECOMMENDATIONS 7-1
REFERENCES R-1

LIST OF ILLUSTRATIONS

Figure | Page

Figure 2-1. Probability Density Function For A Circular Normal Distribution. 2-6
Figure 2-2. Optimal Search Density (Coverage) For A Circular Normal Distribution.............. 2-7
Figure 2-3. Posterior Search Object Location Distribution given Failure to Detect with a Circular Search of Optimally Varying Coverage.. 2-8
Figure 2-4. Detection (vs. Coverage) Functions. 2-9
Figure 2-5. Posterior Search Object Location Distribution given Failure to Detect with a Square Search of Uniform Coverage 2-13
Figure 3-1. Probability Distribution and Search Plan for the *SS Central America*.................... 3-8
Figure 3-2. Probability of Success versus Search Time... 3-8
Figure 4-1. Solid Angle Subtended at the Observer's Eye by a Vessel's Wake........................ 4-7
Figure 4-2. Inverse Cube Lateral Range Curve. 4-9
Figure 4-3. Effect of Five Parallel Tracks. 4-10
Figure 4-4. Inverse Cube POD vs. Coverage (Parallel Sweeps). 4-11
Figure 4-5. "Piled Slab Solid" Approximation to Optimal Coverage. 4-14
Figure 4-6. Parallel Sweep Search.. 4-15
Figure 4-7. Expanding Square Search. 4-15
Figure 4-8. Circular Normal Probability Map. 4-21
Figure 4-9. An Example of a CASP Probability Map... 4-24
Figure 4-10a. Monte Carlo Distribution. 4-26
Figure 4-10b. Monte Carlo Distribution... 4-27
Figure 4-10c. Monte Carlo Distribution. 4-28
Figure 4-10d. Monte Carlo Distribution... 4-29
Figure 4-10e. Monte Carlo Distribution. 4-30
Figure 4-10f. Monte Carlo Distribution.. 4-31
Figure 4-11. Three Detection Functions... 4-33
Figure 4-12. Actual versus Intended Search Craft Track.. 4-36
Figure 4-13. Systematic Relative Motion Distortion. 4-37
Figure 4-14. Un-normalized Update of CASP Probability Map. 4-41
Figure 4-15. Re-normalized Update of CASP Probability Map.. 4-41
Figure 5-1. Min/Max Applied to Leeway Rate. ($D_e = 0.125 \times$ distance drifted)....................... 5-2
Figure 5-2. Min/Max Applied to Leeway Rate. ($D_e = 0.3 \times$ distance drifted).......................... 5-3
Figure 5-3. Min/Max Applied to Leeway Rate. ($DV_e = 0.6$ knots) 5-4

LIST OF ILLUSTRATIONS

Figure — Page

Figure 5-4. Ridge of High Probability Density Between $Datum_{min}$ and $Datum_{max}$. 5-6
Figure 5-5. Min/Max Applied to Leeway Divergence. ($D_e = 0.3 \times$ distance drifted) 5-8
Figure 6-1: SARP Drift Solution 6-6
Figure 6-2. CANSARP Drift Solution. 6-8
Figure 6-3: SARMAP Drift Solution 6-10
Figure 6-4: SARIS Drift Solution 6-12
Figure 6-5: C2PC/AMS Drift Solution 6-16
Figure 6-6. Obtaining a Sample Sea Current Vector. 6-22
Figure 6-7. CASP 1.x Drift Solution 6-24

LIST OF TABLES

Table — Page

Table 3-1. VPCAS (OASIS) Operational Results 3-5
Table 4-1. Koopman Search Radii 4-16
Table 4-2. CSPM Search Radii 4-16
Table 4-3. Comparison of Recommended Search Radii. 4-19

LIST OF ACRONYMS, ABBREVIATIONS AND DEFINITIONS

σ	Standard deviation
σ	Standard error
AMS	Automated Manual Solution
ARCVIEW®	Open-architecture geographic information system (GIS) product
ASA	*Applied Science Associates, Inc.*, a company engaged in marine and fresh water environmental modeling.
ASW	Antisubmarine Warfare
ASWEPS	*U. S. Navy's Antisubmarine Warfare Environmental Prediction System*
ASWORG	U.S. Navy's Antisubmarine Warfare Operations Research Group
BMT	*British Maritime Technology*, a multi-disciplinary engineering and technology consultancy
C2PC	Command and Control Personal Computer
C2PC/AMS	Command and Control Personal Computer/Automated Manual Solution
CANSARP	Canadian Search and Rescue Program
CAS	Computer Assisted Search
CASP	Computer Assisted Search Planning
CDC	Control Data Corporation
COMPATWINGSPAC	U. S. Navy Commander, Patrol Wings Pacific Fleet
COTS	Commercial Off The Shelf
CPA	closest point of approach
CSP	Commence Search Point
CSPM	Classical Search Planning Method
DC	District of Columbia
DMB	Datum Marker Buoy
DoD	Department of Defense
DR	Dead Reckoning
DV_e	probable error of the drift velocity estimate,
E	total probable error of position
Effort	(total length of the searchers' tracks while searching, $L = vt$)
ELT	Emergency Locator Transmitter
EPIRB	Emergency Position Indicating Radio Beacon
EPRB	An ancillary program used with early versions of CASP to assist with the inputting of locations from EPRIBs and ELT's
erf	the error function
FLIR	Forward-looking infrared radar
GDOC/AMM	Geographic Display Operations Computer/Automated Manual Method
GEM®	COTS product from Electronic Information Systems, Inc.
GIS	Geographic Information System
G-OPR	U.S. Coast Guard, Commandant, Office of Search and Rescue
GPS	Global Positioning System
GSO	Generalized Search Optimization
GUI	Graphical User Interface

HM Coastguard	Her Majesty's Coast Guard (United Kingdom)
IAMSAR	*International Aeronautical and Maritime Search and Rescue*
ICAO	International Civil Aviation Organization
ICS	Incident Command System
IMO	International Maritime Organization
IP	initial position
ISOPREP	isolated personnel report
Kts	knots
MEP	Marine Environmental Protection
Min/Max	Minimum / Maximum
NM	nautical miles
NOAA	National Oceanic and Atmospheric Administration
NVG	night vision goggles
NY	New York
OASIS	Operational ASW Search Information System
OEG	U. S. Navy's Operations Evaluation Group (successor to ASWORG)
OILMAP/ARCVIEW®	*ASA Oil spill trajectory and dispersion prediction software using the ARCVIEW® GIS.*
OPS	Ocean Prediction System
^{o}T or T	Degrees True
PLAN	An ancillary program used in conjunction with SARP or CASP in the final stages of search planning. This program divided a large rectangular search area into smaller rectangular sub-areas for assignment to individual SRUs, maintaining a uniform coverage factor throughout.
POC	Probability of Containment
POD	Probability of Detection
POS	Probability of Success
R	radius
R&DC	U. S. Coast Guard Research and Development Center
RCC	Rescue Coordination Center
RDT&E	research, development, test, and evaluation
S	Track Space or Track Spacing
SAR	search and rescue
SARIS	Search and Rescue Information System (developed by *BMT*)
SARLANT	USCG Atlantic Search and Rescue teletype network
SARMAP	Search and Rescue Information System (developed by *ASA*)
SARP	Search and Rescue Planning program
SARPAC	*USCG Pacific Search and Rescue teletype network.*
SARSCENE 99	Canadian National Search and Rescue Conference, 1999
Search Effort	area effectively swept, $Z = W \times L$
SMC	SAR Mission Coordinator
SRU	Search and Rescue Unit
SS	Steam Ship
ENE	East Northeast
TWC	total water current

U.S.	United States
UK CG3	United Kingdom's Coast Guard Search Planning Methodology, version 3
UNIX	a widely used computer operating system
USCG	United States Coast Guard
USS	*United States Ship*
VPCAS	U. S. Navy Aviation Patrol Computer Assisted Search
W	effective sweep width (often shortened to just *sweep width*)
Z	*area effectively swept*

CHAPTER 1.

INTRODUCTION

1.1 BACKGROUND

A prime mission of the U.S. Coast Guard (Coast Guard) is Search and Rescue (SAR). This mission involves rendering assistance to distressed vessels and aircraft and their crews and passengers within the territorial coastal waters of the United States, on certain large lakes (e.g., the U.S. portions of the Great Lakes), and on the high seas in those areas for which the United States has accepted SAR responsibility as a party to an international treaty or convention. Altogether, the Coast Guard's area of SAR responsibility covers more than 28 million square nautical miles. The Coast Guard also cooperates with and, resources permitting, often assists other nations in SAR planning, training and operations by virtue of its position as the world's premier maritime SAR organization. The Coast Guard Research and Development Center (R&DC) supports the marine SAR mission and helps maintain the Coast Guard's position of world leadership with a variety of research, development, test, and evaluation (RDT&E) projects.

For the past 25 years the R&DC has helped the Coast Guard continuously refine, test and evaluate the data, equipment, and procedures used in SAR. The R&DC has addressed a broad spectrum of marine SAR issues covering such topics as improving the ability to predict search object drift, sensor performance testing and evaluation under operational conditions, and evaluation of search planning methods. Much of the emphasis has been placed on acquiring the needed data and developing the methods for using it to determine where and how to search so that the chances of finding the search object are always maximized.

Many SAR events are quite straightforward in nature. A vessel in distress may notify the Coast Guard directly of its situation and precise location. In such a case, the Coast Guard can dispatch or direct the appropriate assets to the scene to affect a rescue. Although this is called a "SAR" incident, a case such as this would be a rescue mission primarily, with little or no searching required.

Although the vast majority of "SAR" cases are straightforward rescue operations, the Coast Guard still must expend significant sums each year on those that are not. The factors that keep these cases from being straightforward rescues usually involve significant uncertainty about the search object's exact whereabouts, adding a significant amount of complexity to the SAR process. Before a rescue can take place, the distressed persons must first be located, and this requires planning and conducting effective, efficient searches. Both of these tasks can be daunting. (Planning and conducting effective, efficient surveillance missions aimed at interdicting illegal drugs and migrants bound for the U.S. are similarly daunting tasks—perhaps even more so since those engaged in illegal activities are often evasive, actively seeking to avoid detection.) For this reason, substantial portions of search and rescue manuals are devoted to the planning and conduct of search operations.

1.2 OVERVIEW

Searching as a general activity has been a common everyday endeavor for millennia. Perhaps this is the reason it has always been generally accepted that searches could consume considerable resources and that success was as much a matter of luck as anything else. Although many searches were no doubt organized affairs, the reasons for organizing them were most likely related to logistics and simple resource management concerns rather than the issues of efficiency and effectiveness. It was not until the demands of modern warfare, anti-submarine warfare in particular, placed such a high premium on the efficiency and effectiveness of search and surveillance operations that they came under serious scientific scrutiny, giving rise to the field of study now known as ***search theory***.

The developments in search theory and its application to SAR may be broken into three categories. These are:

Scientific research into search theory itself and the resulting developments,

Development of search planning doctrine and "how to" methods for SAR manuals, and

Development of computer-based search planning decision support tools.

Unfortunately, these paths have not been as closely coupled as they should have been over the years since World War II when the study of search theory first began. In the chapters that follow, we will look at each of these categories in turn and show where and how the latter two are connected to the first, as well as some instances where doctrine and/or computer models were modified in ways that do not seem to weather scientific scrutiny very well. We will begin with the history of major scientific developments in search theory, starting with the work done during World War II and progressing to the present. This survey of search theory history will include brief looks at some situations where search theory has been successfully applied. Next, we will explain the basic principles of search theory in the simplest possible terms. After that, we will turn the clock back to 1957 when the U.S. Coast Guard first articulated its search planning doctrine in the form of a search and rescue manual and show how the basic principles of search theory were applied to the SAR search problem. We will then follow the developments in that doctrine through the various editions of the *National Search and Rescue Manual* and finally into the recently published *International Aeronautical and Maritime Search and Rescue Manual (IAMSAR Manual).* Finally, we will look at computer-based search planning decision support tools, beginning with the U.S. Coast Guard's *Search and Rescue Planning (SARP)* tool developed circa 1970.

CHAPTER 2.

THE DEVELOPMENT OF SEARCH THEORY

2.1 Initial Development

Search theory is the study of how to most effectively employ limited resources when trying to find an object whose location is not precisely known. The goal is to deploy search assets to maximize the probability of locating the search object with the resources available. Sometimes this goal is stated in terms of minimizing the time to find the search object. Search theorists seek to find methods, procedures, and algorithms that describe how to achieve these goals.

Work on search theory began in the U.S. Navy's Antisubmarine Warfare Operations Research Group (ASWORG) in 1942 in response to the German submarine threat in the Atlantic (see Morse [1982]). A summary of the work done by this group from 1942 to 1945 is given in Sternhell and Thorndike [1946]. Bernard Koopman joined ASWORG in 1943, and at George Kimball's suggestion, Koopman, James Dobbie, and a few others were given the job of assembling the existing results on search into a coherent theory. Morse [1982] credits Koopman with providing the basic probabilistic foundation of the subject and finding the first results on optimal allocation of search effort, specifically the optimal allocation of a fixed amount of search effort using the exponential detection function to detect a stationary search object having a bivariate normal distribution of possible locations.

Koopman defined the elements of the basic problem of optimal search. They are:

A prior probability density distribution on search object location (so the probability of containment (POC) for any subset of the possibility area can be estimated),

A detection function relating search effort density (or coverage, C) and the probability of detecting (POD) the object if it is in a searched area,

A constrained amount of search effort, and

An optimization criterion of maximizing probability of finding the object (probability of success or POS) subject to the constraint on effort.

Finding the allocation (time, place, and amount) over some subset(s) of the possibility area for the limited amount of available search effort that maximizes the probability of success is called the optimal search problem. The solution to this problem tells the search planner the sub-area(s) where search effort should be placed and how much effort should be placed in each.

The resulting synthesis of search theory by Koopman and his colleagues was published in *Search and Screening* (Koopman [1946]) as Operations Evaluation Group (OEG) Report 56. (OEG was a descendant of ASWORG and other research groups where the relatively new applied science of operations research was being applied to naval military problems.) *Search and Screening* defined many of the basic search concepts such as:

Effective search (or sweep) width (W)

Effective search (or sweep) rate (W x search speed)

Lateral range (detection) function (POD as a function of distance off a searcher's single track)

Effort (total length of the searchers' tracks while searching, $L = vt$)

Search Effort (area effectively swept, $Z = W \times L$)

Search Effort Density (coverage, $C = Z$/area searched)

Detection function (i.e., a POD vs. Coverage function, e.g., POD = $1 - e^{-C}$)

Koopman [1946] used these concepts to develop efficient methods for locating stationary objects and also provided methods for designing barrier searches and antisubmarine warfare screens to detect moving adversaries. It presented mathematical models for visual, radar and sonar detection of objects. Koopman [1946] and its updated version, Koopman [1980], are still the classic references on basic search theory.

2.2 TYPES OF SEARCH PROBLEMS

The work of Koopman and his colleagues in the ASWORG (and later the Operations Evaluation Group (OEG)) laid the groundwork for the development of search theory and the applications that followed. It is convenient to categorize this subsequent work according to the type of search problem involved. A detailed bibliography and discussion of the types of search problems can be found in Benkoski *et al.* [1991]. We provide only a brief overview here.

2.2.1 One-Sided Search Problems

The simplest types of search problems are those in which the searcher can choose his strategy, but the search object neither chooses a strategy nor reacts to the search in any way. These are called *one-sided* search problems. Most maritime SAR searches are treated as one-sided search problems. Usually the time and place of the SAR incident and any subsequent movement of the survivors are not deliberately chosen by the persons involved for the purpose of affecting their chances of being found by searchers since they did not plan to be the object of a search. Survivors also often cannot or do not react to searchers in a way that significantly affects the chances for detecting them. The only exception to this rule occurs when the searchers come very close to the survivors and those survivors have signaling devices available and use them in response to the searcher's presence to try to attract the searcher's attention. The simplest one-sided problems involve searching for a stationary search object.

2.2.1.1 Stationary Search Objects

A stationary search object is one that does not move. The searches for the sunken treasure ship, *SS Central America*, the missing submarine *USS Scorpion*, and the H-bomb lost off the coast of Spain in 1966 are examples of searches for stationary search objects. Other examples include searches for downed aircraft, hidden natural resources (gas, oil, minerals, etc.), searches for archeological sites and artifacts, and even searches for something as mundane as lost car keys. These are one-sided search problems because the search object has not chosen its location and it does not react to the searcher's efforts.

2.2.1.2 Moving Search Objects

Search for a life raft adrift in the ocean is an example of a one-sided moving object search problem. The movement of the raft is not (substantially) under the control of the people in the raft, and the people are not able to react to the search effort except perhaps by trying to signal an aircraft or vessel they observe passing nearby. Searches for submarines can be considered one-sided searches when the searching platform or system is covert, i.e., when the target submarine is unaware of the searcher's presence.

2.2.2 Two-sided Search Problems

In two-sided search problems, both the search object and the searcher are allowed to choose their strategies. Two-sided problems can involve either stationary or moving search objects. An example of a two-sided stationary search object problem occurs when the search object chooses a place to hide and stays there. The searcher then has to find the search object. Most two-sided problems involve moving search objects. Two-sided search problems divide into cooperative and non-cooperative searches.

2.2.2.1 Cooperative Searches

An example of a two-sided cooperative search is a rendezvous search. In these searches each party behaves in a manner that maximizes the chances of one party finding the other. Thus each party attempts to make itself as detectable as possible by the other while at the same time making the best use of its own ability to detect the other party. (Sometimes the degree of cooperation is limited by the need for both parties to avoid detection by an adversary, as when searching for a downed allied pilot behind enemy lines.) For example, when two people have become separated in a crowd and wish to find one another again, we have a cooperative search problem. In addition to each person actively looking for the other, further cooperation may take such forms as wearing distinctive clothing or hats, waving, agreeing to be in a limited area within a certain timeframe, etc.

Another example of two-sided search is searching for an intelligent person lost in the woods who understands how the searchers will operate if/when a search effort is mounted. That person may try to move to a place where he can be found more easily or to cooperate in some way by leaving or giving signals to indicate his position. Simply remaining in one place upon realizing one is lost is often a considerable aid to searchers. In many areas, children are taught to "hug a tree"

and stay where they are if they become lost in the woods. (Unfortunately, many children are also taught to avoid strangers, and in more than one case a lost child has delayed his own rescue by failing to respond to search parties passing nearby because they were strangers to him. In fact, there have been times when children (and hunters with delicate egos) have turned what should have been a two-sided cooperative search into a two-sided non-cooperative search—a type discussed briefly below.) Needless to say, two-sided cooperative searches are generally both shorter and more often successful than other types.

2.2.2.2 Non-Cooperative Searches

Many two-sided searches are non-cooperative. One example is one submarine searching for another submarine that is trying to remain undetected when each is aware of the other's presence. Another example is law enforcement officers searching for drug smugglers who are trying to evade detection. Still another example is a manhunt for a suspect or an escaped convict.

2.3 KOOPMAN'S EARLY WORKS

Koopman's original 1946 report, *Search and Screening*, was initially classified. However, Koopman [1956a, 1956b, 1957] published three articles that summarized, in an unclassified fashion, the theoretical aspects of the work reported in *Search and Screening*. In these papers Koopman showed how to find optimal allocations of search effort when the search object is stationary and the detection function is exponential. This included showing how to solve explicitly for the optimal effort allocation for a bivariate normal search object location distribution.

2.3.1 Effective Sweep Width

In his groundbreaking work on search theory, Koopman [1946] defined the "*effective search (or sweep) width*" (often shortened to just *sweep width*) as follows: If a searcher passes through a swarm of identical stationary objects uniformly distributed over a large area, then the *effective search (or sweep) width*, **W**, is defined by the equation,

$$W = \frac{Number\ of\ Objects\ Detected\ Per\ Unit\ Time}{\left(Number\ of\ Objects\ Per\ Unit\ Area\right) \times \left(Searcher\ Speed\right)}, \tag{2-1}$$

where all values are averages over a statistically significant sampling period. If the *lateral range function* is known for a given search situation, then the *area under the lateral range curve* equals the *sweep width, W,* for that situation. That is, if the detection probability is expressed as a function d_r of lateral range x from a sensor's single straight track through the swarm of objects, then

$$W = \int_{-\infty}^{+\infty} d_r(x)dx \tag{2-2}$$

This *effective sweep width* is also twice the maximum detection range of an "equivalent" *definite range detection profile* (one that is 100% effective out to some definite lateral range either side

of its track and completely ineffective beyond that range). Here, "equivalent" means that the definite range detection profile and the actual detection profile both detect, on average, the same number of objects per unit time under the same conditions of object density and searcher speed.

The effective sweep width, W, depends on three classes of factors. These are,

1. The search object's characteristics affecting detection by the sensor(s) in use (object size, color, reflectivity/emission properties, etc.),

2. The capabilities of the sensor(s) in use (visual acuity, a radar's ability to reliably detect a standard test object at various ranges, the effect of speed on performance, etc.), and

3. The environmental conditions at the place and time of the search that affect the performance of the sensor(s) in use (visibility, weather, sea state, vegetation (ground cover), etc.).

2.3.2 Exponential Detection Function

Suppose that we are searching with a sweep width W and moving at speed v. If we are searching uniformly throughout a region of area A with the effectiveness of "random" search, then the probability of detecting the search object by time t given it is located in the region is

$$P(t) = 1 - \exp\left(-\frac{Wvt}{A}\right) \tag{2-3}$$

The fraction Wvt / A is the density of search effort in the region. In operational search and rescue (SAR) terminology, this fraction is called *coverage*. The numerator Wvt is called *search effort* or *area effectively swept* (Z). Thus *coverage* is the ratio of the *area effectively swept* to the amount of *area searched*. Equation [2-3] is called the *random search formula*. This typically gives a lower bound on the effectiveness of a systematic search that tries to spread its effort uniformly over the search region. The term "random search" must not be taken too literally. With completely random searching, one can obtain very non-uniform coverage of the search area, and as a result obtain a lower probability of detection than that given by the "random" search formula.

Suppose that $f(x)$ is the density of search effort (coverage) in the neighborhood of the point x in our search space, the plane. Let

$b(x, f(x,y))$ = probability of detecting the search object given it is located at x,y and the search density is $f(x,y)$.

If

$$b(x, f(x,y)) = 1 - \exp(-f(x,y)), \tag{2-4}$$

we say that b is an *exponential detection (vs. coverage) function*. An exponential detection function means that in each local area, the search has the effectiveness of the random search formula.

2.3.3 Search Object Location Distribution

Suppose that the search object is stationary and our knowledge of the search object's location is given by a bivariate normal probability distribution with its mean (center) at (0, 0). This knowledge may have been obtained from a navigational fix with some uncertainty. The density distribution of possible locations around such a navigational fix is often of the bivariate normal type. Many times it is convenient to take the fix as the origin of our coordinate system so that the mean of the uncertainty distribution is at (0, 0).

The density function, p, for this distribution is given by

$$p(x_1, x_2) = \frac{1}{2\pi\sigma_1\sigma_2} \exp\left[-\frac{1}{2}\left(\frac{x_1^2}{\sigma_1^2} + \frac{x_2^2}{\sigma_2^2}\right)\right] \tag{2-5}$$

A graph of this density function is shown in Figure 2-1 for the case where $\sigma_1 = \sigma_2$. This is called a circular normal distribution. The height of a point on the surface of the "mountain" above the (x_1, x_2) plane represents the probability density at that point. The probability density is highest at the center of the distribution (0,0) and decreases as distance from the center increases. In theory, this distribution covers an infinitely large area since the probability density approaches, but never actually reaches, zero. In practice some reasonable cutoff is applied.

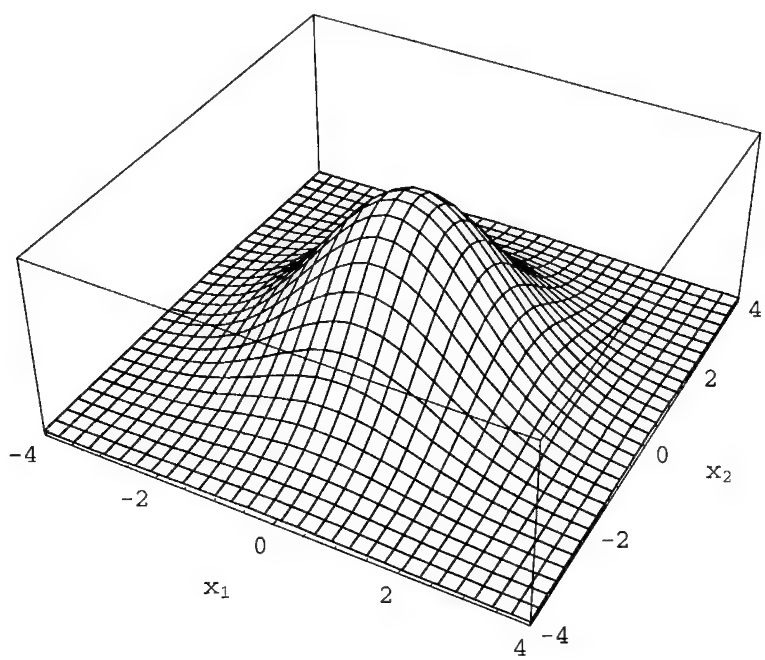

Figure 2-1. Probability Density Function For A Circular Normal Distribution.

2.3.4 Optimal Search Effort Density

Suppose that the sweep width is W when the sensor travels at speed v. If we have T hours of search time available, then Koopman showed that the optimal search effort density (coverage) f^* is

$$f^*(x_1, x_2) = \begin{cases} \left(\dfrac{WvT}{\pi\sigma_1\sigma_2}\right)^{\frac{1}{2}} - \dfrac{1}{2} r^2(x_1, x_2) & \text{for } r^2(x_1, x_2) \le 2\left(\dfrac{WvT}{\pi\sigma_1\sigma_2}\right)^{\frac{1}{2}} \\ 0 & \text{for } r^2(x_1, x_2) > 2\left(\dfrac{WvT}{\pi\sigma_1\sigma_2}\right)^{\frac{1}{2}} \end{cases}$$

where

$$r^2(x_1, x_2) = \frac{x_1^2}{\sigma_1^2} + \frac{x_2^2}{\sigma_2^2}. \tag{2-6}$$

Figure 2-2 shows an example of the optimal search effort density (coverage) for the circular normal location density in Figure 2-1 based on some specific amount of available search effort. Note that the optimal coverage is highest at the center where the location probability density is also highest. The optimal coverage then decreases with distance from the center until at a certain radius where the optimal search effort density becomes zero. In other words, all of the available search effort is expended within a certain radius (depending on the amount of effort available) and none is expended outside that radius even though there is some small probability of the search object being outside the corresponding circle.

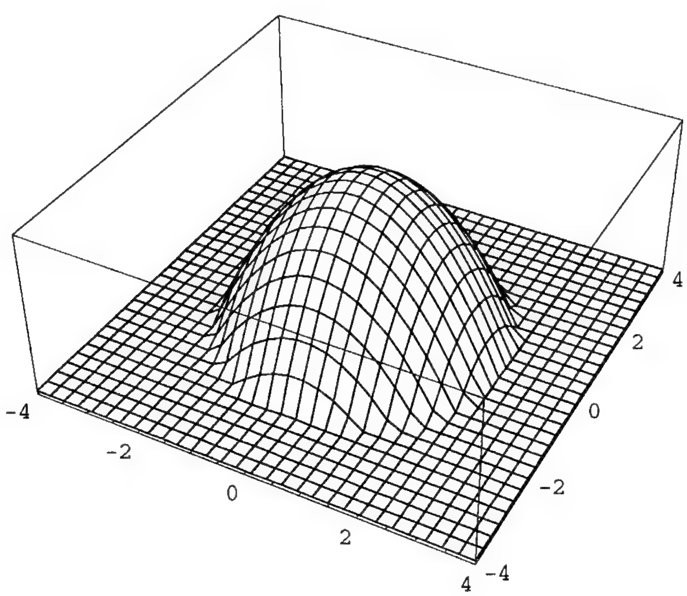

Figure 2-2. Optimal Search Density (Coverage) For A Circular Normal Distribution.

2.3.5 Posterior Search Object Location Density

Suppose that we apply the optimal search density as shown in Figure 2-2 and fail to detect the search object. The resulting distribution is shown in Figure 2-3. This is the posterior search object location density given the search has been unsuccessful.

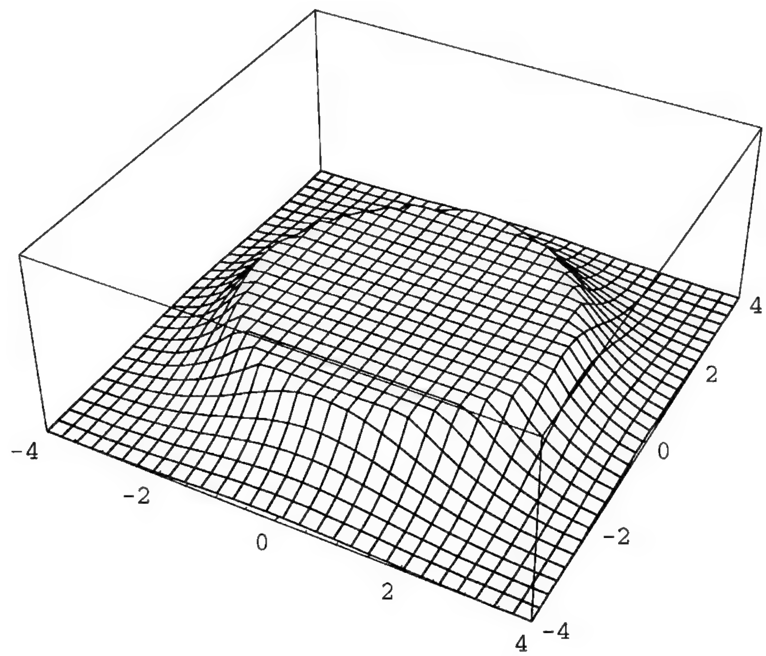

Figure 2-3. Posterior Search Object Location Distribution given Failure to Detect with a Circular Search of Optimally Varying Coverage.

The distribution shown in Figure 2-3 was computed by employing the form of probabilistic reasoning called Bayes' rule. The posterior density is flat inside the circle where search effort has been applied. As more and more effort is applied (in an optimal fashion), the height of this "mesa" (posterior density) becomes lower and the radius of the circle of search increases making the distribution as a whole flatter. Koopman [1957] later extended his optimal allocation results from normal distributions to a more general class of probability distributions.

2.4 Non-Exponential Detection Functions and Cellular Distributions

Koopman's results have been extended in two important directions. The first allows us to find optimal allocations of search effort when the detection function is not exponential and the second deals with search object location distributions composed of independent, possibly non-contiguous, cells. Some of the standard detection models used by the U.S. Coast Guard, such as the inverse-cube model initially postulated by Koopman [1946] as a model for visual detection, produce non-exponential detection functions. Cellular search object location distributions commonly arise in both land and maritime SAR situations.

2.4.1 Non-exponential Detection Functions

Koopman studied two other detection (vs. coverage) functions besides the so-called random search function. These were the "definite range" and "inverse cube" detection functions for parallel track search patterns. All three detection functions are shown in Figure 2-4.

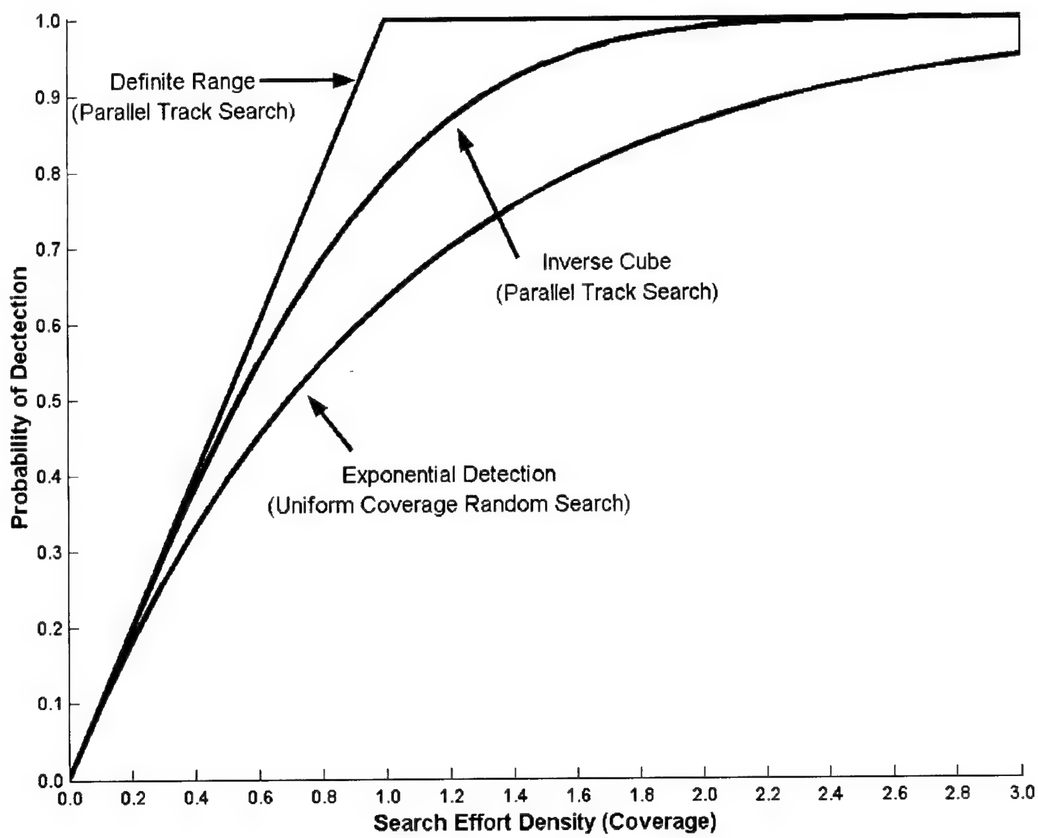

Figure 2-4. Detection (vs. Coverage) Functions.

In the case of definite range detection, detection of the search object is guaranteed to occur whenever the distance between the sensor and the object is less than or equal to the sensor's "definite detection range" for the object in question under the environmental conditions at the time and place of the search. There is never any detection beyond this range. Therefore, the probability of detection at any instant is purely binary—one or zero—depending on the distance separating the object and the sensor. When employed in parallel track searching, the probability of detection increases linearly with the search effort density (coverage) until both reach 1.0. Higher effort densities can produce no better result so the detection probabilities remain constant at 1.0 for all densities (coverages) greater than one. Although the definite range detection function is not very realistic, it does provide an upper bound on all detection functions and has some other mathematical uses as well.

The inverse cube detection (vs. coverage) function was the result of Koopman's efforts to model visual detection performance mathematically. Koopman reasoned that the instantaneous (or "one glimpse") probability of detecting an object visually is proportional to the solid angle subtended by the object at the observer's eye. In working out the geometric and mathematical consequences of this assumption, Koopman found that the instantaneous probability of detecting the object was inversely proportional to the cube of the distance between the observer and the object. This relationship is the origin of the name "inverse cube law of visual detection." Koopman went on to show that in a parallel track search pattern, this model of visual detection produces a detection function given by

$$P(t) = erf\left(\frac{Wvt\sqrt{\pi}}{2A}\right) \text{ or } P(C) = erf\left(\frac{\sqrt{\pi}}{2}C\right) \tag{2-7}$$

The exponential and inverse cube detection functions have an important property. They exhibit a *decreasing rate of return.* This means that the probability of detection increases more and more slowly as the search effort density (coverage) increases. This effect is seen clearly in Figure 2-4. In mathematical terms, this property is expressed by saying the detection function has a decreasing derivative. A decreasing rate of return is a common property in economic situations in which effort may be measured in dollars, time, or manpower and return is in dollars. Most detection functions have the decreasing rate of return property.

DeGuenin [1961] extended Koopman's results by finding the optimal allocation of search effort for a stationary search object for any detection function with a decreasing rate of return. Detection functions that include the origin (0,0) and have a continuous, positive and strictly decreasing first derivative (rate of return) are called *regular detection functions.*

2.4.2 Optimal Allocations for Cellular Distributions

In the above examples, the probability distributions have density functions that vary smoothly over space. By contrast, Charnes and Cooper [1958] considered situations in which the search space is divided into cells with

p_j = probability the search object is in cell j
A_j = area of cell j
W_j = sweep width in cell j
v_j = search speed in cell j
t_j = time spent searching in cell j.

The search problem is to divide the total search time T over the cells to maximize probability of success. For an exponential detection function, Charnes and Cooper presented an algorithm for computing the optimal distribution of searching effort over these cells to maximize probability of success by search time T.

The problem solved by Charnes and Cooper arises often in land search and rescue. The probability distribution for the location of the search object, say a lost boy, is often cellular with

cells of varying size. Because of variations in terrain, both the sweep width and the speed at which ground parties can search may vary from segment to segment. If used, the Charnes-Cooper algorithm could tell the search planner how many resource (searcher) hours out of the total number available should be devoted to each segment of the search area in order to have the greatest probability of a successful outcome.

2.4.3 Uniformly Optimal Search Plans

The search plans described above maximize the probability of detecting the search object by time T. They tell us the total effort to put into each cell or region, but they say nothing about how the effort should be put into cells over time. Suppose that we want a search plan that tells us how to allocate search effort in both space and time so that at each time t between 0 and T, we have done as well as possible. In fact, we would like the result after any time t of searching to be *optimal* for time t. A plan with this pleasing property is called *uniformly optimal.* Uniformly optimal plans also minimize the mean time to find the search object.

Koopman showed that when the detection function is exponential, a uniformly optimal plan exists. In fact, the way to obtain a uniformly optimal plan is to organize the search effort so that by time t you have allocated search effort to be optimal for that time. One then continues on to a plan that is optimal for time $T > t$ by adding the additional effort in each cell that is required by the time T plan over the time t plan. Stone [1989] showed that uniformly optimal plans exist and may be constructed in a similar fashion whenever the detection function has the decreasing-rate-of-return property.

2.5 OPTIMAL SEARCH FOR MOVING AND MULTI-STATE SEARCH OBJECTS

2.5.1 Optimal Search for Moving Search Objects

Prior to Brown [1980], optimal allocation results for moving search objects were limited to very special cases. Most moving search object problems were approached by freezing the search object motion over some increment of time, allocating effort as though the search object were stationary during that time increment, and then repeating the process for the next time increment. The U.S. Coast Guard's Computer Assisted Search Planning (CASP) System, discussed in Richardson and Discenza [1980], still uses this technique by applying an adaptation of the Charnes-Cooper algorithm to the search object location probability density distribution at a specific point in time, usually the planned commence search time.

Brown discovered an efficient algorithm for finding optimal search allocations for moving search object problems when the search object motion is Markovian and the detection function is exponential. Brown's algorithm maximizes detection probability at time T. The U.S. Navy in searching for Soviet submarines applied this algorithm to great effect. This application is discussed further below.

Washburn [1983] generalized Brown's algorithm to the class of forward and backward (FAB) algorithms that apply to a more general class of payoff functions. Algorithms for non-exponential detection functions and non-Markovian motions are given in Stone [1979] and Stromquist and Stone [1981].

2.5.2 Optimal Survivor Search

In the search problems discussed above, the goal is to maximize the probability of detecting the search object by some time. In the case of search and rescue problems, a more appropriate goal may be to maximize the probability of detecting the search object alive. A search with this goal may apply the initial effort in some lower probability areas that are particularly hazardous in order to recover a survivor quickly if he or she is located there. This may involve some sacrifice of overall probability of success. As an example, one might want to concentrate initially on search areas where a survivor would be located if he is immersed in the water and somewhat delay searching areas that would be likely only if he or she is still in a disabled boat.

Discenza and Stone [1981] developed algorithms for solving this (moving search object) problem and for solving a more general class of problems called multi-state search problems. A multi-state state search problem is one that involves searching for objects that may undergo changes in state following the initial distress incident, such as a person in a disabled boat abandoning it for a life raft and then possibly becoming a person in the water.

2.6 Constraints on the Searcher

In the search problems considered above, we have assumed that effort can be distributed over the search space any way that we choose. Sometimes this is a reasonable approximation. A visual search by aircraft over a geographically limited region where the time to travel from one part of the region to another is small is an example. Sometimes the constraints on the movement of the search platforms require that we consider special types of search plans. Usually there are two of types of constraints that are considered—path constraints and simplicity constraints.

2.6.1 Path Constraints

If the search platform is a boat or a person walking on land, or even an aircraft with a large assigned search area, then the place where the platform is searching now significantly constrains the places where it can search in the next increment of time. In these cases, we have an *optimal searcher path problem.* Instead of finding an optimal allocation of search effort, the problem is to find an optimal path for the searcher. The set of paths from which the optimum is chosen is restricted to those that obey the physical constraints on the movement of the search platform. This is a difficult class of problems, especially for moving search objects, but there has been some progress in solving them. Stewart [1979, 1980], Eagle [1984], and Eagle and Yee [1990] have applied integer programming approaches to finding efficient algorithms for solving these problems. Optimal searcher path problems are basically equivalent to the well-known NP-complete "traveling salesman" problem.

2.6.2 Simplicity Constraints

In executing actual searches, it may be desirable to restrict the search patterns to a class of searches that are simple to execute operationally. A typical example is to restrict search plans to be composed of searches consisting of a set of rectangles each with a uniform search density or coverage. Such plans can be approximated by searches that employ straight, equally spaced,

parallel search paths in the rectangles. In fact, this is the method usually employed when searching large areas from aircraft.

Single Rectangle Searches. In the case of search for a stationary search object with a bivariate normal location distribution, Richardson and Discenza [1980] show how to find optimal rectangle plans. For such a plan, one chooses a single rectangle and spreads the search effort uniformly over that rectangle as illustrated by Figure 2-5. Richardson and Discenza show that it is always possible to pick an optimal rectangle (uniform coverage) plan that comes within 3% of the probability of success produced by Koopman's optimal (non-uniform coverage) plan.

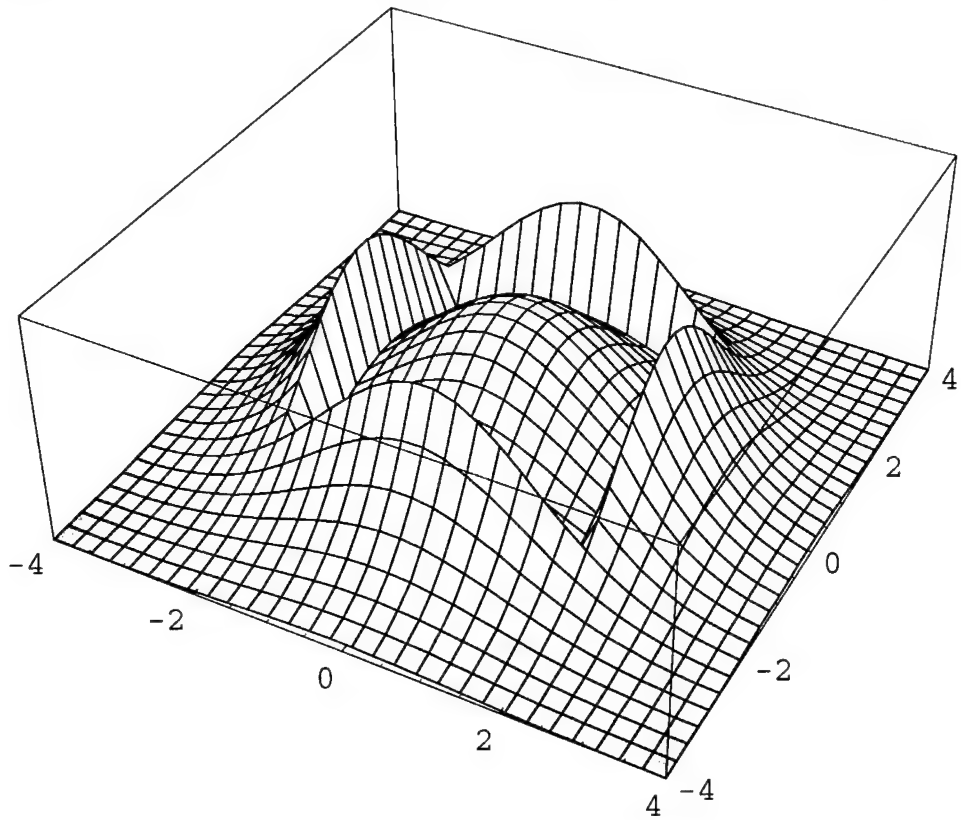

Figure 2-5. Posterior Search Object Location Distribution given Failure to Detect with a Square Search of Uniform Coverage.

Multiple Rectangle Searches. Discenza [1980] developed an algorithm for finding optimal multiple rectangle searches for the cellular search object location distributions generated by CASP. These multiple rectangle searches consist of non-overlapping rectangles. In each rectangle the search effort is spread uniformly over the rectangle. Furthermore, each search asset (say an aircraft) is assigned to search one and only one rectangle. The solution method proposed by Discenza involves some additional restrictions on the choices of rectangles to allow an efficient solution of this problem.

2.7 SEARCH AND EVASION PROBLEMS

A classic two-sided search problem involves a search object that is trying to evade a searcher. In one case the search object's goal may be simply to avoid detection. In other cases, the search object may have additional goals such as reaching a certain area undetected. This would be the goal for a smuggler or an infiltrator. The problem for the search theorist to solve is finding the optimal strategy for both the searcher and the evader. These problems tend to have a game theory formulation. Although they are difficult to solve, there has been some progress made by Auger [1991], Eagle and Washburn [1991], Gal [1980], Stewart [1981], and Washburn [1980].

CHAPTER 3.

PRACTICAL APPLICATIONS OF SEARCH THEORY

3.1 RELATING THEORY AND PRACTICE

The previous chapter emphasized the results of basic research into the theory of search. This chapter will look at some examples of how the principles of search theory have been applied to obtain practical solutions to "real world" problems.

As a prelude to discussing practical solutions to "real world" problems, Koopman [1980] points out several important facts. The first is that the number and variety of constraints imposed on an actual search operation, "...while perfectly obvious to anyone engaged in carrying out the operation, are of so varied and irregular a nature that they defy precise mathematical formulation." In other words, the constraints discussed in the previous chapter are but a small subset of those with which a search planner must deal and there is no mathematical way to factor all the additional constraints into the optimal effort allocation problem. One consequence is that even when a theoretically optimal allocation of effort can be computed, only a very crude approximation of that allocation can actually be realized with the available resources. However, as Koopman also points out, "There is nothing peculiar to the theory of search in this; it affects every practical implementation of the theoretical results of operations research."

Koopman goes on to describe how theoretical results relate to practical operations:

> "In spite of such discordance between what can be shown to be optimal and what can actually be done, the theoretical developments make, among others, three essential contributions to practical matters. First, they show directions in which the practicable operations should be changed in order to improve them. Second, by calculating a theoretical optimum or measure of effectiveness and then gauging the calculated or observed results of a practicable operation against such measures, they can indicate whether the [results] are unreasonably poor and ought to be improved. Third, they can serve a useful purpose in *limiting arguments*, often valuable in excluding basically fruitless operations."

Perhaps the most valuable contribution of theoretical results is providing a good starting point for planning efficient, effective search, surveillance and screening operations. For example, it is often possible to provide planners with methods for easily computing an optimal allocation of effort subject to only a few basic constraints and assumptions. Suppose a planner computes an optimal allocation subject to a constraint on the amount of effort available, a requirement to search with uniform coverage and a known or assumed detection (vs. coverage) function. This allocation ignores many practical physical constraints of the problem, such as those governing the physically possible movements of the sensor platform. However, given the answer to the less constrained effort allocation problem, the planner can then use his knowledge of what is and is not operationally feasible to design an operational plan that approximates the less constrained optimal plan as closely as the myriad additional practical constraints will allow. Plans developed

in this fashion will almost always be superior to those based only on "hunches," intuition, past experience, or "judgment."

There is another more subtle, but no less important, benefit to using theory-based approaches to search planning. In the editor's preface to Stone [1989], John D. Kettelle makes the following observations:

> "The search process in inherently a nervous one. Either you will find the 'target' or you won't. This involves more stress than the continuous penalties or payoffs associated with dullness or brilliance in dealing with problems such as scheduling or logistics. This discontinuity makes search a little like litigation. During an actual 'case' there is a sense of urgency and emergency. This stress can trigger a major, sometimes frantic, effort. Experts can be mobilized. Armies (or navies) can be sent scurrying around. A nervous principal or client can make intuitive decisions that are painfully wrong."

An orderly planning approach based on proven theories minimizes the likelihood for making "intuitive decisions that are painfully wrong." It also minimizes the likelihood of wasting scarce resources on ill-conceived plans and truly random, ineffective "helter-skelter" searches.

3.2 NAVAL APPLICATIONS DURING WORLD WAR II

In the last three chapters of *Search and Screening*, Koopman [1980] describes several ways in which search theory was successfully applied to specific types of naval warfare problems during the Second World War. These included searching for targets in transit, setting up sonar screens to protect convoys and task forces from submarine attack, and aerial escort of convoys and task forces, again to protect against submarine attack. Of the three, the type that is most germane to U.S. Coast Guard peacetime operations is searching for craft in transit.

Koopman [1980] repeats his detailed description and discussion of how to construct "crossover" barrier patrols, initially given in Koopman [1946]. He also illustrates why such a technique is needed by showing the ineffectiveness of simple parallel sweeps across the adversary's intended track due to unswept areas relative to the search object. The objective of a barrier patrol is to prevent the undetected transit of an adversary through a "channel" (which may be some region of open ocean). It is generally assumed that the search object's *intent* and *capabilities* are reasonably well known. To know the intent is to know the adversary's approximate route, i.e., what geographical area it is coming from, approximately where it is going and areas it intends to transit to get from one to the other. To know the adversary's capabilities is to know such things as the range of possible speeds, endurance at the various speeds, etc. The essential things to have are good estimates of the adversary's course and speed in an area that meets two requirements: (a) the adversary intends to pass through it, and (b) the area can be effectively covered by one's own forces (i.e., it cannot be too far from one's own staging area(s).) The barrier is considered effective when it detects adversaries attempting to cross it (assuming that such detection always makes it possible to successfully thwart the adversary). It is also considered effective when the adversary abandons his objective to avoid detection.

Koopman provides three practical sample problems to illustrate how "crossover" barriers were successfully applied to intercepting blockade-runners during World War II. He goes on to note that this type of search is feasible only when the searcher's speed exceeds that of the search object by a substantial margin, such as in the case of patrol aircraft forming a barrier for enemy ships. Barriers for situations where such disparities in speed do not exist are possible, but depend on so many factors that vary from case to case that there is no practical general solution. Each must be solved individually.

U.S. Coast Guard law enforcement patrols to prevent the illegal entry of drugs, other types of contraband, and migrants are considered effective if they detect those heading for U.S. shores for illegal purposes, assuming detection implies successful interception and arrest. Law enforcement patrols are also considered effective if smugglers and illegal migrants cease attempting to achieve their objectives due to fear of detection. However, the overall effort may not be considered effective over time if it merely forces smugglers to shift their operations from one part of the coast to another that is less effectively patrolled.

For search and rescue, one might be tempted to regard drifting survivors as analogous to "targets in transit" since they move from one general location to another with "courses" and "speeds" that can be predicted, if only approximately. Unfortunately, barrier searches are of only limited usefulness in SAR. A barrier search is effective against only a small range of search object speeds near the average speed for which the barrier is designed. Although drift rates are slow, it is quite possible to have a large percentage error between the predicted and actual values. In addition, drifting search objects do not tend to move along long, straight paths but instead exhibit only trends with a significant degree of randomness in their movements. Finally, in many places where a barrier might seem most appropriate, such as the Straits of Florida or across a river or estuary, it is known that there will be a wide variation in search object speeds along any barrier that is established. A barrier that is effective for small objects (e.g., life rafts) near the center of the Florida Straits moving northward at four knots will be largely ineffective for similar objects nearer the edges moving northward at less than two knots.

Koopman [1946, 1980] also describes an expanding square search designed to prevent the undetected movement of a search object away from an area about a point of fix to another location (the so-called "fleeing datum" problem). However, he first describes how to develop a near-optimal uniform coverage expanding square search plan for a stationary object having a circular bivariate normal distribution of locations. Koopman provides a formula for computing the inscribed radius, s_k, of each near-optimal search square in a series of n searches:

$$s_k = \frac{\sigma}{2}\sqrt{\pi \frac{W}{S}(2k-1)}, \quad k = 1, 2, \ldots, n \tag{3-1}$$

The squares computed by this method define the bases of rectangular prisms with height W/S which, when stacked on one another to form a "stepped pyramid," approximate the paraboloid of revolution in Figure 2-2 that represents optimal coverage. We will return to this formula when we compare the "classical search planning method" from early versions of the Coast Guard and National SAR Manuals to formal search theory.

3.3 More Recent Naval Applications

The U.S. Navy has used theory-based search planning techniques extensively in many, if not most, of its search problems—including many classified applications. Computerized implementations are now the order of the day and are often integrated with signal processing and navigation systems.

3.3.1 Finding Lost Objects

The Navy has conducted numerous searches for lost objects. Two that received a lot of publicity are the search for the H-bomb lost off the Mediterranean coast of Spain in 1966 and the search for the wreck of the submarine *USS Scorpion* in 1968.

H-bomb Search. In the H-bomb search, Richardson [1967] used computers and a multiple scenario approach to develop the probability map for the bomb's location. Because of the limitations of computers in 1966, the on-scene computation of the effects of unsuccessful search had to be performed manually, by Richardson. Fortunately, this was a search for a stationary search object so a manual procedure was reasonable.

Scorpion Search. The same principles were applied to the *Scorpion* search. Multiple scenarios were developed for the location of the *Scorpion*. The uncertainties in each of the scenarios were quantified through the use of probabilities, and a probability distribution for the location of *Scorpion* was produced by Monte Carlo simulation on a computer. Analysts were sent on-scene to update the distribution for search effort and to recommend the allocation of the continuing search effort. This work is documented in Richardson and Stone [1971]. Again this process was strikingly successful. The Scorpion was found within 260 yards of the highest probability cell in the distribution.

In both of the above cases, successful outcomes would have been extremely unlikely without the use of techniques based on search theory.

3.3.2 Soviet Submarines

In the early 1970s the ideas behind the Computer Assisted Search Planning (CASP) system being developed for the U.S. Coast Guard were used by H. R. Richardson to develop computer search planning programs for the U.S. Navy. These programs were initially used to help plan searches for Soviet submarines in the Atlantic. This was later extended to the Mediterranean and the Pacific. The most complete version of this system was developed for use in planning searches by antisubmarine warfare patrol aircraft for Soviet submarines patrolling in the Pacific in the late 1970s. This system was called OASIS and later VPCAS.

VPCAS. VPCAS contained a number of improvements and extensions over the CASP system. Namely,

VPCAS was capable of using historical information to form the probabilistic models used for submarine motion.

As well as accounting for the effect of unsuccessful search though the use of detailed models of sonobuoy detection capability, VPCAS incorporated detection information from underwater surveillance sensors.

VPCAS used the optimal-allocation-for-moving-targets algorithm developed by Brown [1980] to recommend the locations of sequences of sonobuoy fields over several days to maximize probability of success.

Success of VPCAS. For a certain period of time, as this system was introduced to the personnel who planned searches, some searches were planned using the standard manual techniques and some were planned using VPCAS (OASIS). At the end of this period, analysts compared the effectiveness of searches that were planned with the computer system to those that were planned manually and discovered a striking result. *The probability of success for searches planned with VPCAS was twice as high as the probability of success for searches planned manually.* Table 3-1 presents an unclassified version of the results of the analyses performed by Benkoski [1978] and McCoy [1978]. The numbers in parentheses indicate the number of searches over which the percentage was calculated. The cases indicate variations on the definition of success, but all the definitions involve detection.

Table 3-1. VPCAS (OASIS) Operational Results.

CASE	**VPCAS Success %**	**Manual Success %**
I-A	73% (48)	32% (157)
I-B	56% (48)	20% (157)
II-A	82% (17)	43% (65)
II-B	65% (17)	25% (65)
III-A	71% (17)	32% (65)
III-B	53% (17)	17% (65)

Upon looking at these results, one might wonder if the computer system was used on the easy searches thereby giving it an unfair advantage. Benkoski and McCoy looked at that possibility and found exactly the opposite. VPCAS tended to be used on the harder searches because it was less familiar than the manual system. Operators would use it only when they felt they needed it, and that was on the more difficult searches.

3.3.3 Clearance

In 1974, the U.S. Navy assisted the Egyptians in clearing unexploded ordnance from the Suez Canal following the settlement of the 1973 Yom Kippur war between Egypt and Israel. Computer search planning and evaluation of the clearance effort played a major role in assuring

that a thorough clearance job was done. Search planning played a role in a number of aspects of the clearance operation.

Computers were used to develop plans for testing and estimating the effectiveness of the search sensors, primarily side-looking sonars. The search effectiveness of the sensors varied as the conditions in the canal varied. Test plans were generated for each area. Using the results of the tests, lateral range curves were estimated by statistical inference techniques.

Using the performance estimates generated in-situ, computer systems were used to generate search plans that produced a high probability of detecting all unexploded ordnance.

Computers were used to cluster the contacts produced by the side-looking sonar and to develop areas for divers to search in order to locate and identify likely contacts.

As the operation proceeded, careful estimates were made of the actual clearance percentage that was obtained. If the after-the-fact estimate was too low, additional clearance effort was applied to improve the clearance probability.

3.4 HISTORICAL WRECKS—*SS CENTRAL AMERICA*

Search theory was used to plan the search for the famous historical wreck of the *SS Central America.* In 1857, while carrying passengers and gold from California to New York, the *SS Central America* sank in a hurricane nearly 200 miles from land, taking gold bars and coins worth an estimated 400 million dollars to the ocean bottom almost 8000 feet below. Some 425 people, including the captain, lost their lives. In 1989, after only three short summer sorties to the area, the Columbus-America Discovery Group had located the wreck and recovered one ton of gold bars and coins from it. This is in sharp contrast to typical treasure hunting operations where individuals spend many years, or even entire lifetimes, in unsuccessful efforts involving far less difficult search conditions.

In 1985, Stone [1992] was given the task of developing a probability distribution for the location of the *Central America.* This distribution was used for constructing the search plan that found the wreck. The methods used to develop the distribution were based on classical search theory techniques and included a combination of historical, statistical, analytic, and subjective methods. This wreck had been the object of many previous searches, but none used a systematic search theory approach and none were successful.

The first part of the search problem was to estimate the more likely, and less likely, locations for the wreck of the *Central America.* Following the paradigm developed by Koopman [1946] and his colleagues in the Navy's Operations Evaluation Group during World War II, this estimate was produced in the form of a two dimensional probability distribution on the location of the wreck. To develop this distribution, Stone made use of the following information:

Historical:

- Herndon's last reported position as passed to the schooner *El Dorado,*
- Sighting of the *Central America* by the brig *Marine,*
- Recovery of survivors by the bark *Ellen,*
- An estimate of the wreck's location by Captain Badger, a passenger on the *Central America,*
- The drift of the survivors on the raft,
- Estimates of wind speed and direction recorded during the hurricane;

Statistical:

- Assumed probability density distributions around reported/estimated positions
- Historical distribution of winds and currents in the area;

Analytical:

- Estimates of the uncertainty in celestial navigation,
- Estimates of the effect of wind on the drift of the *Central America,*
- Estimates of the wind-driven current;

Subjective:

- Weights representing the quality of the information used to estimate the wreck's location.

The methodology for combining these diverse types of information had its start with work done by Richardson [1967] during the 1966 search for the H-bomb lost off the Mediterranean coast of Spain. In many search problems the information about the search object's location comes from a variety of sources and is often inconsistent. The information does tend to cluster into self-consistent sets, however, each of which tells a "story" leading to possible locations of the search object. These clusters are called scenarios. Because of the inconsistencies among the scenarios, one cannot combine them in a standard statistical fashion as though they were independent and unbiased estimates of the search object's location.

In 1968, this search methodology was further developed by Richardson and Stone [1971] to produce probability maps for the successful search for the remains of the nuclear submarine, *USS Scorpion.* The technology reached a more advanced state of maturity in the CASP developed for the U.S. Coast Guard by Richardson and Discenza [1980] with assistance from Stone and others. In fact, Stone [1992] used a modified version of the CASP software to produce probability maps for the location of the *Central America* based on each scenario. The resulting combined distribution was used, in conjunction with estimates of sensor performance developed by Newton [1986], to produce an efficient search plan that yielded a high probability of success. The resulting probability map and search plan are shown in Figure 3-1. This plan consists of a series of long straight legs spiraling out from the high probability areas of the distribution. Long legs were desirable because the sensor, a side-looking sonar, was to be towed at the end of very long cable. A graph of probability of success as a function of search time is shown in Figure 3-2.

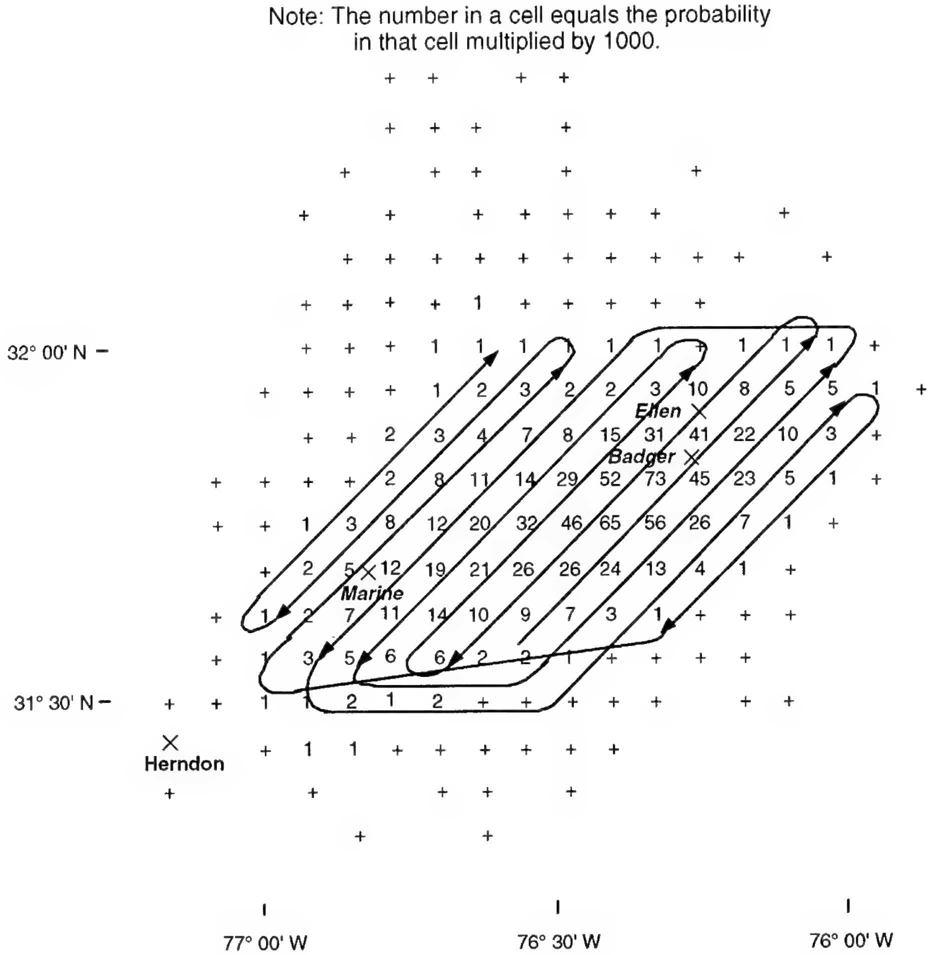

Figure 3-1. Probability Distribution and Search Plan for the *SS Central America.*

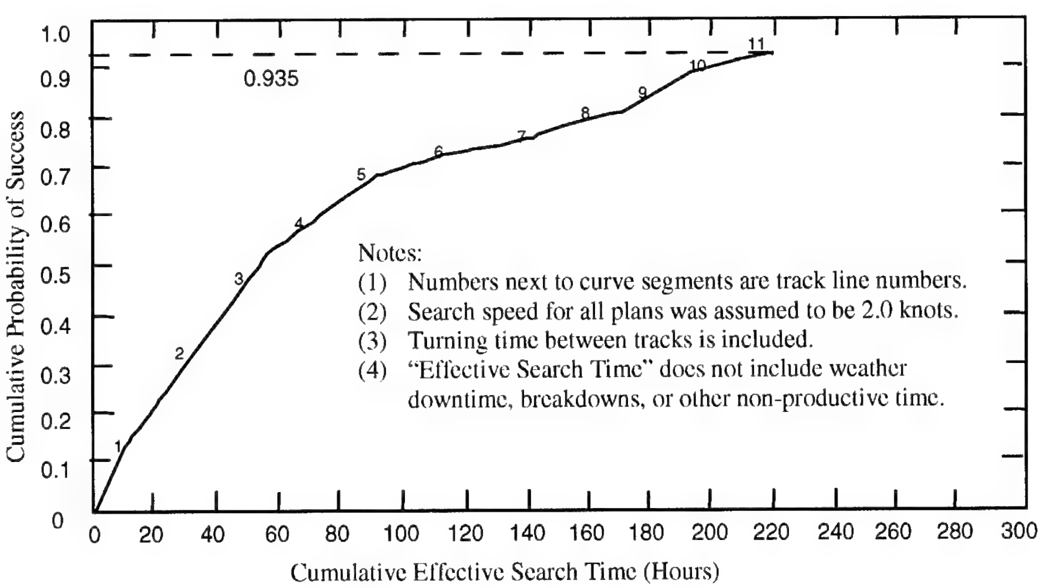

Figure 3-2. Probability of Success versus Search Time.

This successful search provides another example of the effectiveness of systematic searches that are planned using the concepts and methods that have been developed in search theory.

3.5 THE EFFICACY OF THEORY-BASED SEARCH PLANNING TECHNIQUES

The examples given above show that search plans based on search theory often succeed where less scientific methods fail. It is abundantly clear that more scientific search planning using the levels of computing power now commonly available at low cost offers significant benefits. Although implementing scientific methods is not without its costs, these costs are quickly recovered by increased effectiveness in terms of mission performance (e.g., lives saved, smugglers interdicted, etc.) and decreased average time to locate search objects. Decreasing the mean time to detect search objects saves resource hours and that translates directly into monetary savings. Since both the search and surveillance missions of the Coast Guard depend heavily on aircraft, which are expensive to operate, such savings add up quickly.

CHAPTER 4.

CLASSICAL SEARCH PLANNING METHOD

4.1 INTRODUCTION

In 1957 the U.S. Coast Guard published its first search and rescue manual. This manual was to become the basis for the *National Search and Rescue Manual* that replaced it just two years later when the *National Search and Rescue Plan* was first adopted. Like its successors ever since, a large portion of this first manual was concerned with search planning. As we shall see, the search planning methodology contained in this manual was clearly based on the earlier work of Koopman. The fact that the first SAR manual was published just after Koopman's [1956a, 1956b, 1957] unclassified articles appeared is unlikely to have been mere coincidence. Unfortunately, the original work that transformed the theory developed by Koopman into a practical method for planning SAR searches seems to have been lost. Thus if we are to understand the specific connections between search theory and search planning doctrine, we are now faced with the somewhat daunting task of reconstructing the original developer's work as best we can.

We shall proceed through the basic elements of the optimal search problem as given originally by Koopman. These elements are:

A prior probability density distribution on search object location (so the probability of containment, POC, for any subset of the possibility area can be estimated),

A detection function relating search effort density (or coverage, C) and the probability of detecting (POD) the object if it is in a searched area,

A constrained amount of search effort, and

An optimization criterion of maximizing probability of finding the object (probability of success or POS) subject to the constraint on effort.

As we investigate each element, it will be well to bear in mind that the classical search planning method (CSPM) had to be kept as simple as possible. All planning had to be done with nothing more than pencil, paper, standard navigational tools (dividers, parallel rules, etc.), and paper charts. Computations had to be simple enough to be done quickly by hand or with the aid of graphs, nomograms, etc. These requirements meant that some gross simplifications and sweeping generalizations had to be made. We will point these out as we go along.

Another important point to remember is that the classical search planning method is a complete system where each element depends on the nature and form of all the other elements, as well as on their quantitative measures. For example, a distribution of effort that is optimal for one type of search object location probability density distribution will not generally be optimal for any other type.

4.2 PRIOR PROBABILITY DENSITY DISTRIBUTIONS

The first and most obvious assumption is that the search object is located somewhere on the surface of the earth. Since the fraction of the earth's surface that contains all possible search object locations is typically quite small, the area of interest may be satisfactorily approximated by a two-dimensional plane surface.

4.2.1 Incident Position

The simplest SAR situation is when the distressed craft is able to provide its position at the time of the distress. For example, a sinking vessel might issue a distress call giving its position and notifying recipients that the crew was abandoning ship. Very often the incident position is established by a navigational fix or is estimated from a recent fix. In 1957 the state of navigation technology made it common for such positions to be somewhat in error. Position errors were often normally distributed with respect to parallels and meridians, giving the possible positions a bivariate normal distribution with the reported position as the mean. The simplest bivariate normal distribution is the circular bivariate normal distribution where the standard deviations of the normal distributions along the meridians and parallels are equal and the two distributions are completely independent of one another (i.e., uncorrelated). The probability that the actual position is within radius *R* (expressed as a multiple of the standard deviation, σ) of the mean (center) is given by

$$POC = 1 - e^{-\frac{R^2}{2}}$$

Although the standard deviation (or "standard error") is the usual method for quantifying the amount of error, or "spread," of a distribution about the mean value, the quantity used for characterizing position error is *probable error*. For a bivariate normal distribution, the probable error is the elliptical (or circular) contour centered on the mean that contains 50 percent of the distribution. For a circular normal distribution, the relationship between probable error, ε, and the standard deviation, σ, is given by

$$\varepsilon = \sigma\sqrt{-2\ln(1-0.5)} \approx 1.18\sigma$$

If the search object were known to be stationary (e.g., an aircraft that made a forced landing ashore and reported its position immediately prior to reaching the ground), then the assumed prior distribution would be a circular bivariate normal distribution with an estimated probable error (usually based on the navigational capability of the distressed craft). Koopman [1946, 1980] shows how to produce optimal search plans for this type of prior distribution both with and without the constraint of uniform coverage.

4.2.2 Post-incident Search Object Motion and Datum

Very often search planners must contend with post-incident motion. On land, for example, survivors may wander away from an aircraft's forced landing site. We shall look at the marine

environment where objects on the ocean's surface tend to drift away from the incident position under the influence of wind and current.

If there is any significant time lag between a distress incident at sea and the arrival of SAR resources on scene, the survivors will have drifted away from the incident position. In this situation, a search planner will have to estimate a new mean location, or datum, for the time at which search facilities will be searching. The search planner will also have to increase the probable position error estimate to account for the uncertainties added by not having exact values for the wind and current and not knowing exactly how drifting objects respond to these forces.

These added uncertainties raise a potential problem. If the distribution of position errors resulting from uncertainties about search object drift is not circular normal, then the distribution about the new "drifted" datum will not be circular normal either, even if the distribution of errors about the incident position was circular normal. Such added complexity would mean that Koopman's solutions for optimal effort allocation over circular normal distributions could not be used and would also make search planning too difficult for manual implementation. To take advantage of Koopman's work and keep the problem tractable for manual calculation, the CSPM assumed that the distribution of position errors resulting from drift uncertainty was circular normal. This assumption, together with the assumption of a single point datum, requires a number of other implicit assumptions.

If a single drift trajectory is to characterize the search object's motion except for random errors, then it must be assumed that the drift takes place in a homogeneous environment. That is, for all possible search object locations at the beginning of a drift interval, the same distribution of possible drift vectors applies. In other words, neither the mean drift nor the probable drift error is allowed to vary from one place to another, at least not for the duration of a drift interval between datums, i.e., between the incident time and the first search or between sequential searches in a series.

A mean drift vector, or more precisely a mean drift distance vector, is computed as the product of the mean drift velocity and the amount of elapsed time since the incident time or time of the last datum. The mean drift velocity is the vector sum of the mean leeway and mean total water current velocity vectors. The total water current in turn is the vector sum of contributing currents (sea current, local wind current, etc.). If the distribution of errors about the mean drift distance vector is to be circular normal, then the distribution of errors about the mean drift velocity vector must also be circular normal. This means the distributions of errors about the mean leeway and all the mean water current vectors must be circular normal as well.

4.2.3 Applicable Theorems from Statistics

We have now discussed necessary and sufficient conditions to assure a circular normal distribution of possible search object locations about the new datum position. However, we have not provided methods for computing the new datum position or quantifying the probable error about that datum position. To do this, we will need some basic theorems from statistics.

The theorems we need are those that deal with the distribution that results from adding two or more independent distributions together and the distribution that results from multiplying a

distribution by a constant. We add leeway and current distributions to obtain the distribution of drift velocities. We then multiply this distribution by the elapsed time, Δt, to get the distribution of drift displacements. Finally, we add the distribution of drift displacements to the distribution of possible incident positions to obtain the distribution of possible locations for the first search. The following theorems apply to this problem.

Let Y denote a distribution of random variates, let Y_i denote the ith distribution of random variates in a set of mutually independent such distributions, let $\mu\{\}$ denote the mean (expected value) of the distribution contained in the braces { }, and let $\varepsilon\{\}$ denote the probable error of the distribution contained in the braces { }.

Theorem 1: The mean of the sum of any number of independent probability density distributions equals the sum of the means of the independent distributions. That is,

$$\mu\left\{\sum_{i=1}^{n} Y_i\right\} = \sum_{i=1}^{n} \mu\{Y_i\}$$

Theorem 2. The mean of a constant multiple of a distribution equals the product of the constant and the mean of the original distribution. That is,

$$\mu\{cY\} = c\mu\{Y\}$$

Theorem 3: The probable error of the sum of any number of independent probability density distributions is the square root of the sum of the squared probable errors of the independent distributions. That is,

$$\varepsilon\left\{\sum_{i-1}^{n} Y_i\right\} = \sqrt{\sum_{i=1}^{n} \varepsilon^2\{Y_i\}}.$$

Theorem 4: The probable error of a constant scalar multiple of a distribution is equal to the product of the scalar constant and the probable error of the distribution. That is,

$$\varepsilon\{cY\} = c\varepsilon\{Y\}.$$

We will now show how these theorems are used to compute a datum distribution from an earlier incident distribution.

4.2.4 Computing Datum

We will now lay out a skeletal description of the steps used to arrive at an oceanic datum. All speeds are in knots, all distances are in nautical miles, and all times are in hours. All referenced graphs were present in the SAR manual.

1. Estimate the mean incident position and the probable error (X) of the estimate.

2. Consult an appropriate data source to obtain an estimate of the mean wind (direction and speed) for the area surrounding the incident position over the period of time extending from the incident time to the time of the first search.

3. Using the mean wind speed, enter the wind current graph to get an estimate of the mean wind current speed. Add (or subtract) the appropriate angle off the downwind direction to obtain the mean wind current direction.

4. Consult an appropriate data source to obtain an estimate of the mean sea current (speed and direction) for the area surrounding the incident position.

5. Add the mean wind current and mean sea current estimates in vector fashion to obtain the mean total water current. (*Theorem 1* applies.)

6. Using the mean wind speed, enter the leeway graph to obtain an estimate of the mean leeway speed. Assume the leeway is in the downwind direction.

7. Add the mean leeway to the mean total water current in vector fashion to get the mean drift velocity (speed and direction). (*Theorem 1*)

8. Multiply the mean drift velocity by the amount of time between the incident and the first search to obtain the mean drift displacement (direction and distance). (*Theorem 2*)

9. Add the mean drift displacement vector to the mean incident position to estimate the mean datum position. (*Theorem 1*)

10. Estimate the probable error of the mean drift displacement vector as one-eighth of its magnitude. (D_e = 1/8 of the distance between incident and datum positions.)

11. Estimate the probable position error of the search facility (Y).

12. Compute the total probable error in the mean datum position relative to the search facility using the formula (*Theorem 3*)

$$E = \sqrt{X^2 + D_e^2 + Y^2}$$

13. Use E with appropriate "safety factors" (discussed later) to determine search area size.

Historical note: The formula contained in the original 1957 USCG SAR Manual was incorrectly stated as

$$c = d_e + \sqrt{A^2 + B^2}$$

where there was an exact definitional correspondence between c and E, d_e and D_e, A and X, and B and Y. Amendment 3 to the National SAR Manual corrected the error in 1963.

The distribution of possible search object locations implied by the above computational steps is centered on the datum position, is circular normal, and has a probable error of *E*.

> **Assumption 1:** The classical search planning method (CSPM) assumes that the distribution of search object location probability density is defined by a single mean position (known as datum) that has a circular normal distribution of possible errors characterized by a known or estimated radius of probable error.

4.3 RELATING *POD* AND EFFORT DENSITY (COVERAGE)

During the Second World War when Koopman began his work, electronic sensing technology was under intensive study and development. A natural product of the development process was a great deal of detailed information about how these new sensors performed. Such information significantly aided in the application of search theory to the development of effective tactics for using these sensors. However, the reliability of electronic sensors in these early years was relatively low and it was not possible to install them in aircraft until relatively late in the war. Visual search from aircraft remained in widespread use and is still the primary technique used in SAR.

Unlike the new electronic sensors under development, there were no data available on the performance of aircrews searching visually. Due to wartime pressures there were neither time nor resources available to conduct any studies of visual detection. Therefore, Koopman developed a hypothetical mathematical model of visual detection. He then used this model to relate *POD* to search effort density (coverage) for search patterns employing long, straight, parallel, equally spaced tracks relative to the search object to approximate uniform coverage of an area.

4.3.1 Koopman's "Inverse Cube" Model of Visual Detection

Koopman began by making certain assumptions about the nature of the Navy's search problem in WW II. These assumptions were:

1. The observer (searcher) is in an aircraft flying at some height *h* above the ocean's surface.

2. The search object (target) is a vessel cruising on the surface of the ocean.

3. The observer initially detects a cruising vessel by seeing its wake.

4. The instantaneous (one-glimpse) probability of detecting a cruising vessel is proportional to the solid angle subtended at the observer's eye by the vessel's wake.

4.3.1.1 Instantaneous Visual Detection Probability

Figure 4-1 illustrates the last assumption. The observer is at point O and the rectangle on the surface with dimensions a and b represents the vessel's wake.

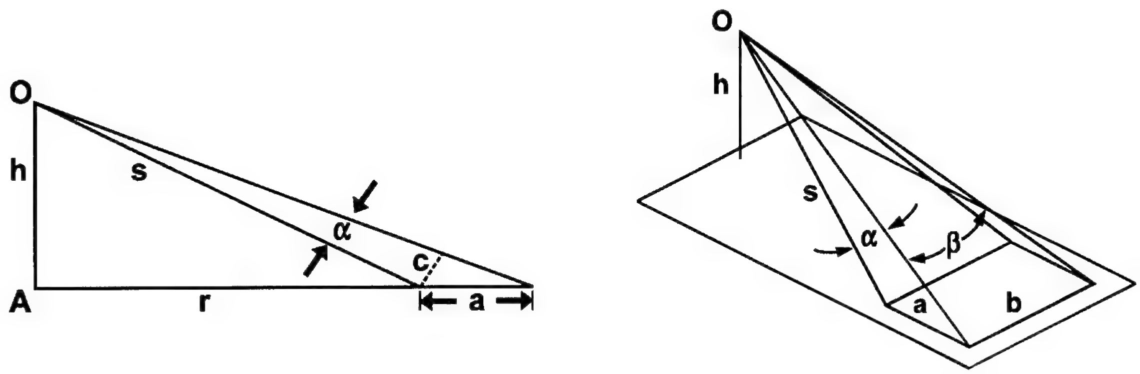

Figure 4-1. Solid Angle Subtended at the Observer's Eye by a Vessel's Wake.

Working through the geometry and associated mathematics, Koopman shows that when a and b are small in comparison to h, r, and s, the solid angle defined by the product $\alpha\beta$ (in radians) is given by

$$\alpha\beta = \frac{abh}{s^3} = \frac{abh}{\left(h^2 + r^2\right)^{3/2}}$$

Since the instantaneous probability, γ, of detecting the wake is assumed to be proportional to the solid angle, then

$$\gamma = \frac{kh}{s^3} = \frac{kh}{\left(h^2 + r^2\right)^{3/2}} \approx \frac{kh}{r^3}$$

The final approximation in this sequence was considered valid because h was small compared to s and r in the majority of cases. The constant of proportionality, k, contains the wake's dimensions, or more precisely its area ab. Stated another way, one of the factors on which k depends is the search object's size. In addition, this constant also depends on all other factors affecting detection that are considered to have fixed values during the search. For SAR, a list of such additional factors is likely to include other search object characteristics (e.g., color), environmental conditions (visibility, weather, etc.), number of observers in the aircraft and their fields of view with respect to the surface, crew fatigue, and many others. In fact, the value of k depends on the same three classes of factors as the sweep width (see paragraph 2.3.1). The relation

$$\gamma \approx \frac{kh}{r^3} \tag{4-1}$$

states that the instantaneous probability of detection is inversely proportional to the cube of the range from the observer to the search object. Hence, Koopman called this model the *inverse cube law* of visual detection.

4.3.1.2 *Lateral Range Detection Function*

Koopman's next step was to investigate the dependence of detection probabilities on the sensor's track relative to the search object. A case of particular interest to the Navy was one where both the sensor and the search object maintain constant speeds and courses over a considerable period of time. In this case, the sensor's track relative to the search object (or the search object's track relative to the sensor) is a long straight line. In this instance, the lateral range is simply the range at the closest point of approach (CPA). If we make the observer the origin of our coordinate system and orient the y-axis so it is parallel to the search object's relative track, we may compute the *detection potential* for an object at lateral range x by accumulating the instantaneous detection probabilities as the object approaches, passes, and moves away from the sensor in relative space. This involves computing the *line integral* along the object's relative track. Using the simplified inverse cube relationship of [4-1] and assuming an infinitely long relative track, Koopman showed that the detection probability as a function of lateral range x is given by

$$d_r(x)=1-e^{-\frac{2m}{x^2}} \tag{4-2}$$

where

$$m=\frac{kh}{v}$$

where v is the speed of the object relative to the sensor. Koopman goes on to derive another relationship for the simplified inverse cube law of visual detection by integrating the right side of equation [4-2] to find the area under the inverse cube lateral range curve. The result is,

$$W = 2\sqrt{2\pi m} \tag{4-3}$$

where W is the effective sweep width. Solving [4-3] for m yields

$$m=\frac{W^2}{8\pi} \tag{4-4}$$

and substituting [4-4] into [4-2] produces

$$d(x)=1-e^{-\frac{1}{4\pi}\left(\frac{W}{x}\right)^2}. \tag{4-5}$$

Graphing this function produces the lateral range curve shown in Figure [4-2].

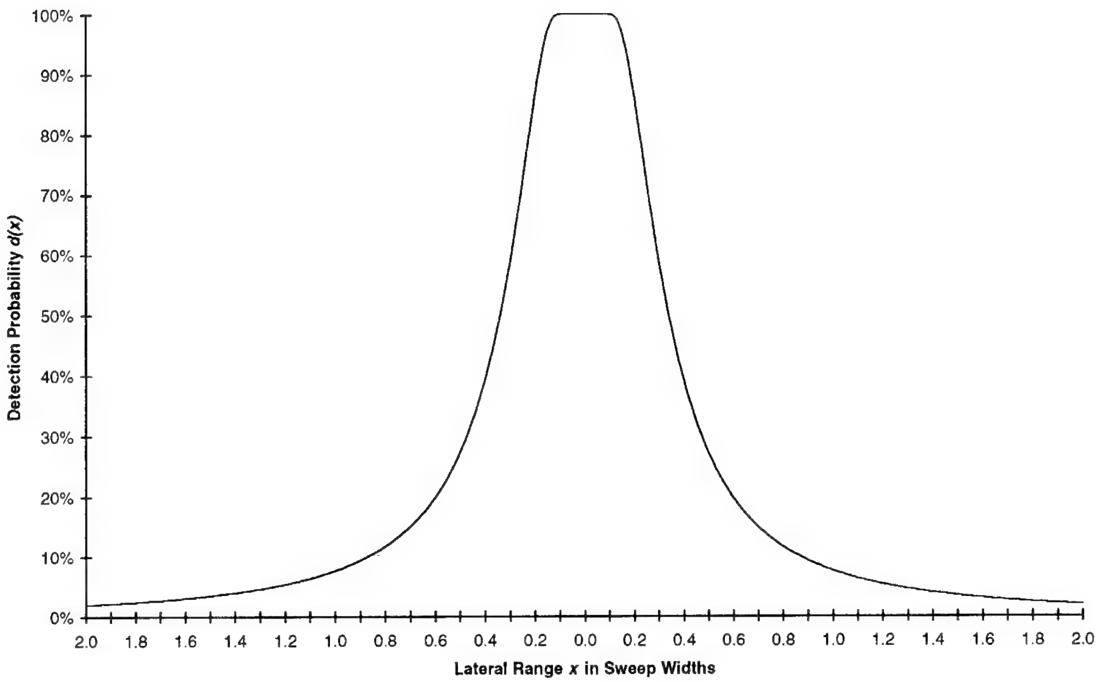

Figure 4-2. Inverse Cube Lateral Range Curve.

Note that in the case of the simplified inverse cube relationship [4-1], the lateral range curve attains a maximum height of 100 percent (guaranteed detection) on the sensor's track and, in theory, the maximum detection range is infinite.

4.3.1.3 The POD vs. Effort Density (Coverage) Detection Function

The final step in Koopman's analysis of his hypothetical inverse cube visual detection model was determining the relationship between the density of searching effort in an area and the probability of detecting the search object if it was in the area during the search. Since a common method for conducting search operations in an area was to move the sensor along a series of long, straight, equally spaced parallel tracks, Koopman evaluated the effectiveness of using an inverse cube (instantaneous) detection function in this fashion. The tracks were assumed to have all the characteristics just listed ***relative to the search object***.

One way to visualize the effect of such a plan is to construct a graph of adjacent lateral range curves and graph their cumulative effective detection probabilities as shown in Figure 4-3.

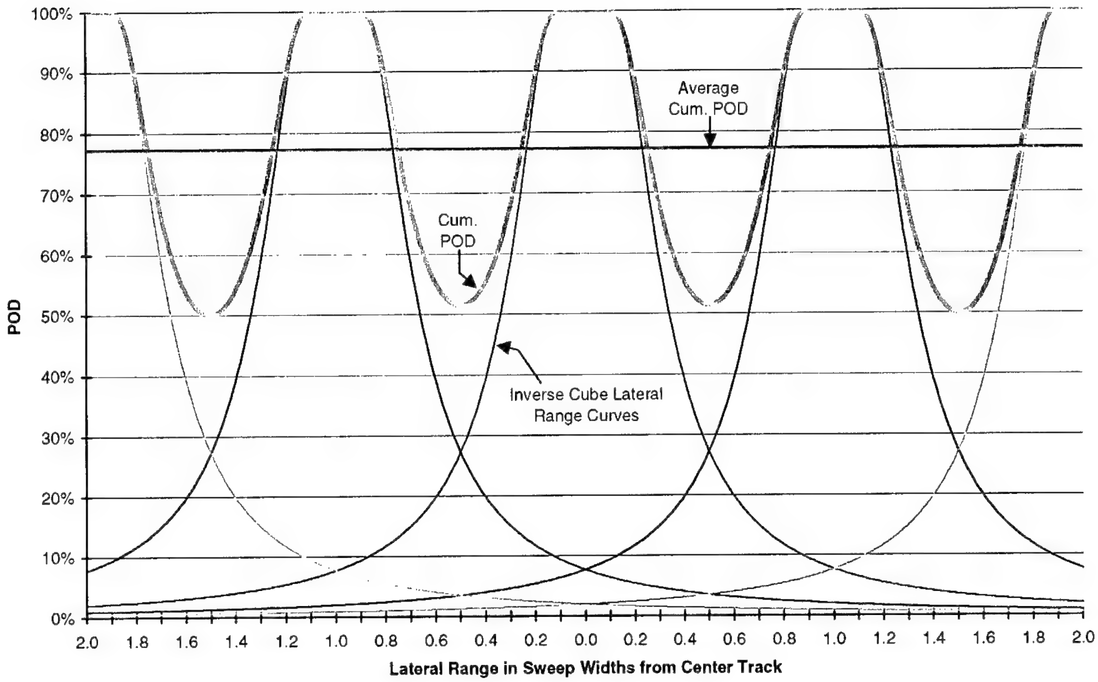

Figure 4-3. Effect of Five Parallel Tracks.

In Figure 4-3, the spacing between tracks is equal to the effective sweep width. For parallel track search patterns covering rectangular search areas, this produces an average search effort density (coverage) of 1.0. For the five tracks shown in Figure 4-3, the average cumulative detection probability across the four sweep widths shown is about 77.3 percent. The traditionally accepted value for a search at a coverage of 1.0 is 78 percent. Koopman showed that for many infinitely long straight parallel tracks spaced one sweep width apart, the average cumulative POD should be about 79 percent. In fact, Koopman derived the general POD vs. Coverage detection function for this type of search using the simplified inverse cube model of visual detection and found that

$$POD = erf\left(\frac{\sqrt{\pi}}{2}\frac{W}{S}\right) \tag{4-6}$$

where *erf* is the well-known error function

$$erf(x) = \frac{2}{\sqrt{\pi}}\int_0^x e^{-t^2}\,dt\,,$$

W is the effective sweep width and *S* is the distance between adjacent parallel tracks (i.e., track spacing). The graph of this detection (vs. coverage = *W/S*) function is shown in Figure 4-4. This is the POD vs. Coverage curve that has appeared in the *National SAR Manual* for many years.

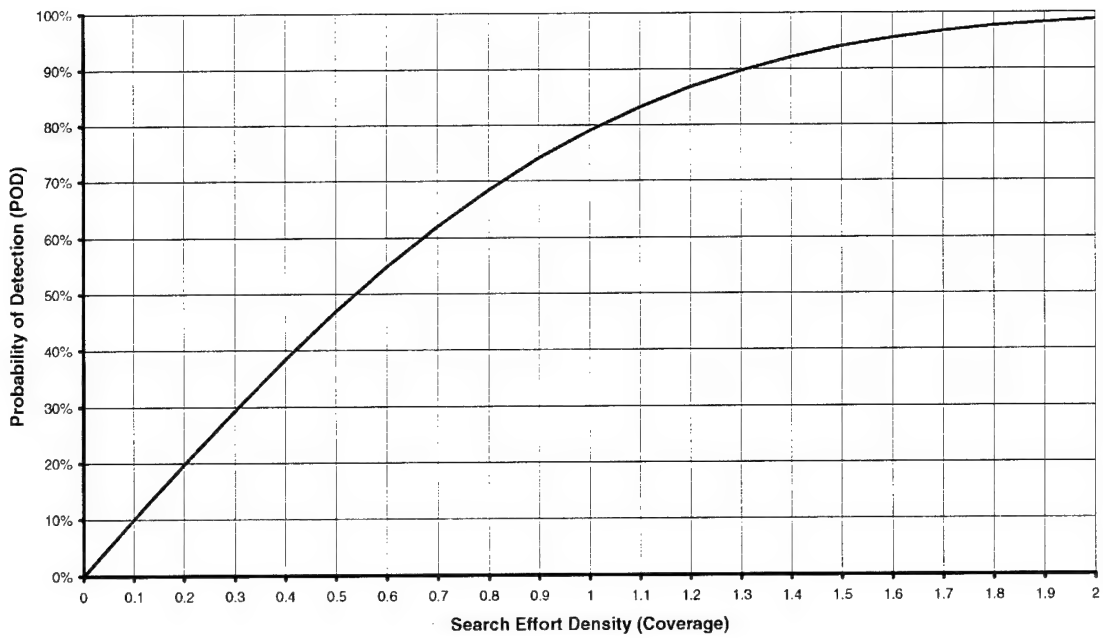

Figure 4-4. Inverse Cube POD vs. Coverage (Parallel Sweeps).

Historical note: The POD vs. Coverage graph contained in the original 1957 USCG SAR Manual was that of erf(W/S) and not that of Equation [4-6]. Amendment 3 to the National SAR Manual corrected this error in 1963.

Assumption 2: The classical search planning method (CSPM) assumes that the detection (vs. coverage) function is based upon the inverse cube law of visual detection being applied under uniform search conditions by using search patterns consisting of long, straight, equally spaced, parallel tracks relative to the search object.

4.4 CONSTRAINTS ON EFFORT

A major factor affecting every search plan is the constraint on the amount of effort available for searching. This constraint is imposed by "real-world" limitations of various sorts. For example, there are no "perfect" sensors that can guarantee detection with the expenditure of a finite amount of effort, even if the area containing the search object is limited. The number of available search platforms is always limited. In addition, these platforms have limited ranges and they and their crews have limited times of endurance. Although all these constraints are real and important, the CSPM did not define *effort* or *search effort*. Nevertheless, it was clear even in 1957 that the size of the recommended search area needed to be related in some way to the amount of available search effort. In the paragraphs that follow, we will define *effort* and *search effort.* In the next section we will examine the CSPM's recommendations for sizing and covering search areas and will compare these recommendations to optimal effort allocations.

4.4.1 Constraints on *Effort* and *Search Effort*

Effort, as used here, is defined as the distance, L, traversed by a search platform while searching in a defined geographical area. If the average speed, v, of the platform while searching and the time, t, spent searching are known, then the effort, z, may be computed as

$$z = L = vt\,. \tag{4-7}$$

Search Effort, Z, is defined as the amount of *area effectively swept* by expending effort in a defined search area. The amount of area effectively swept is the product of the effective sweep width and the distance traversed by the search platform while searching in the area. That is,

$$Z = Wz = Wvt\,. \tag{4-8}$$

Before proceeding to discuss how the CSPM dealt with the constraint on available search effort, it may be well to examine some relationships among the three variables on which it depends.

The search speed, v, depends on the capabilities of the search platform. The minimum and maximum possible speeds clearly depend upon the type of platform in use. However, the search speed affects both of the other two variables—sweep width and time on scene. Generally speaking, increasing speed decreases time on scene because fuel is being consumed at a faster rate. However, distance is also being covered at a faster rate. For most search facilities, there is usually some "most economical" speed that maximizes the product of v and t, i.e., maximizes the distance that can be traversed with the available fuel. However, the search speed can also affect the effective sweep width, W. This is particularly true for aircraft. Usually, increasing speed decreases sweep width. For searching, the optimum search speed would be, in theory, the one that maximized the area effectively swept with the available fuel or in the available time.

Time on scene is not necessarily limited by the type of platform in use or its fuel capacity. For standard visual search, the amount of available daylight may be the limiting factor. The approach of inclement weather may limit the time on scene, either by substantially reducing sweep width or by creating conditions too dangerous for the search facilities. The endurance of the search facility may exceed that of its crew, making crew endurance the limiting factor rather than platform endurance. In some environments, the expected or maximum survival time of the victims of a SAR incident may also limit the available time on scene if the survivors are to be saved.

Sweep width is an obvious constraint on the amount of area that can be effectively swept. It is also a quantity over which the search planner has little or no control.

4.4.2 Other Constraints on *Effort* and *Search Effort*

In addition to the above constraints, we also have the path and simplicity constraints described in Chapter 2. The path constraint arises because real search facilities cannot instantly move from one place to another far away. They must follow some track through space and time subject to their physical limitations (speed, rate of turn, minimum turning radius, etc.). The simplicity constraint ultimately arises from path constraints but there is also a practical need to specify

search plans in simple forms that are easy to understand and carry out. Simplicity also aids flight safety.

Assumption 2 of the last section largely takes the path and simplicity constraints into account. We will simply add the observation here that search patterns meeting the requirements of Assumption 2 are, operationally, rectangular in shape. Adding the symmetry of the circular normal distribution of possible search object locations described in Assumption 1 to the mix produces a square search area centered on datum.

4.4.3 Accounting for Constraints on *Effort* and *Search Effort*

Knowing the limitations on the available search effort is important for the same reasons as knowing the resource limitations for the accomplishment of any task. The available resources need to be applied in a manner that maximizes the chances for successful completion of the task. When planning a search, the objective is to maximize the probability of success by applying the available resources in an optimal manner subject to the various "real-world" constraints. Ideally, a search planning methodology would provide a means for the search planner to obtain the appropriate information about the available resources, use this information to estimate the amount of available search effort at the scene, and produce a workable search plan that maximized POS. Based on Assumptions 1 and 2 above, the problem reduces to that of finding the optimal size for a square search area centered on datum as a function of available search effort. In reality this is a very complex and difficult thing to do, especially for searches subsequent to the first search. This complexity is no doubt why the developers of the CSPM did not provide any means for determining search area size from the amount of available effort and the total probable error of the datum position. Instead, these developers took a different tack, described in the next section.

> **Assumption 3:** The classical search planning method (CSPM) assumes that the optimal search area is a square centered on datum that is searched with a uniform density of search effort (uniform coverage).

4.5 OPTIMAL EFFORT ALLOCATION

All of the assumptions described so far are either obvious from a direct inspection of the CSPM (e.g., square search areas centered on datum) or are easily shown to be derived directly from either elementary statistics (e.g., the formula for estimating total probable error of position) or Koopman's work (e.g., the POD vs. Coverage graph). However, the rationale behind the CSPM's technique for determining search area size is not immediately apparent.

4.5.1 Koopman's Uniform Coverage Squares

Koopman [1946, 1980] described a method for determining the sizes of a series of uniform coverage square search areas over a circular normal distribution of search object location probability density. The technique was based on the following line of reasoning: Approximate the optimal coverage given by the paraboloid of revolution in Figure 2-2 with a series of successively larger uniform coverage square searches so the effort density forms a "piled slab solid," like that depicted in Figure 4-5, having the same volume as the paraboloid of revolution.

In both cases, the enclosed volume represents the total amount of available effort. In the case of the paraboloid this effort was applied in a continuously varying density—a technique that is operationally infeasible. The "piled slab solid," on the other hand, divides the available effort into a number of unequal discrete amounts represented by the volumes of the individual slabs. When the effort is applied in these amounts from smallest to largest (top to bottom in Figure 4-5), an operationally feasible approximation to the ideal variable coverage solution results. Koopman also reasoned that since the search object location probability density was highest at the center, the accumulated effort density (coverage) at the center should always be optimal. Put another way, the height of the "stepped pyramid" of successive uniform coverage square searches should equal the height of the paraboloid of revolution being approximated.

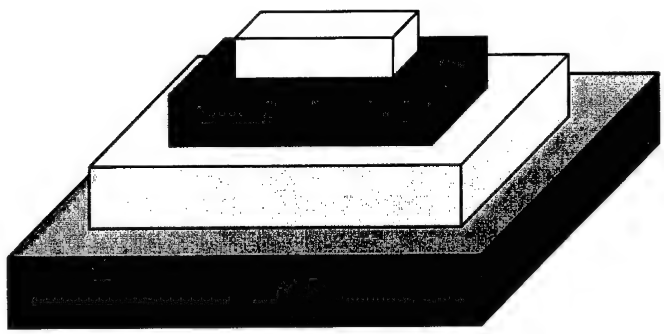

Figure 4-5. "Piled Slab Solid" Approximation to Optimal Coverage.

In Figure 4-5, the height of each "slab" or layer equals the effort density or coverage and the square base of each slab represents the search area. Koopman derived a formula, equation [3-1], for computing the "optimal" search radius (inscribed radius of the "optimal" search square) for each search in a series of searches where the coverages for all searches in the series are equal. Converting equation [3-1] to use the total probable error of the datum position, E, instead of the standard error, σ, we get,

$$s_k = \frac{E}{2.355}\sqrt{\pi \frac{W}{S}(2k-1)}, \quad k = 1, 2, \ldots, n. \qquad (4\text{-}9)$$

For search patterns that employ straight equally spaced parallel tracks, the ratio W/S gives the coverage, C, for a rectangular area whose sides extend one-half of the track spacing beyond the search and cross legs of the pattern. In Figure 4-6, the ratio W/S correctly computes the coverage or average effort density within the dashed rectangle.

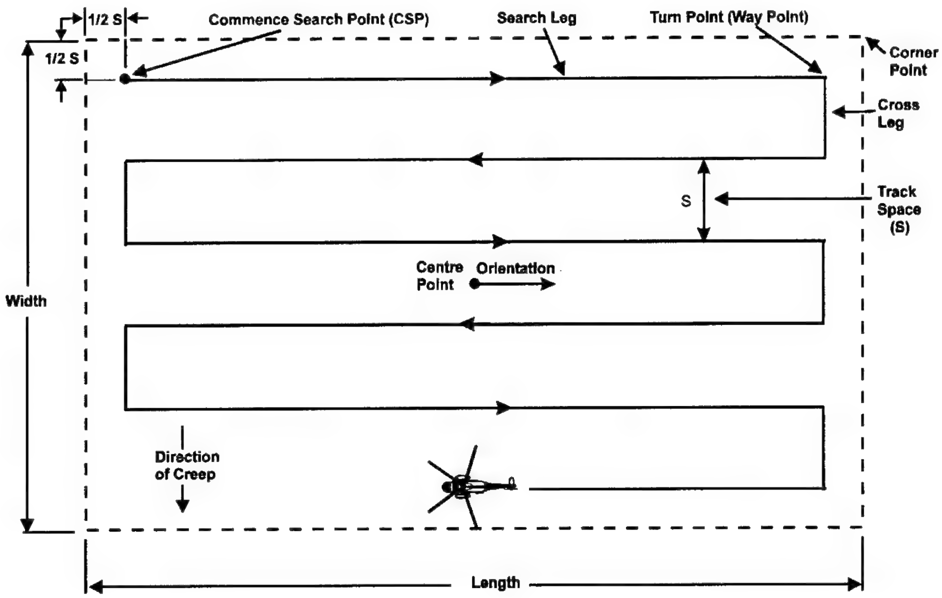

Figure 4-6. Parallel Sweep Search.

Koopman provided no general guidance regarding what effort density (coverage) should be used for the "slabs" in Figure 4-5. However, he did have a suggestion for the case where these slabs were searched by means of an expanding square search pattern like that depicted in Figure 4-7.

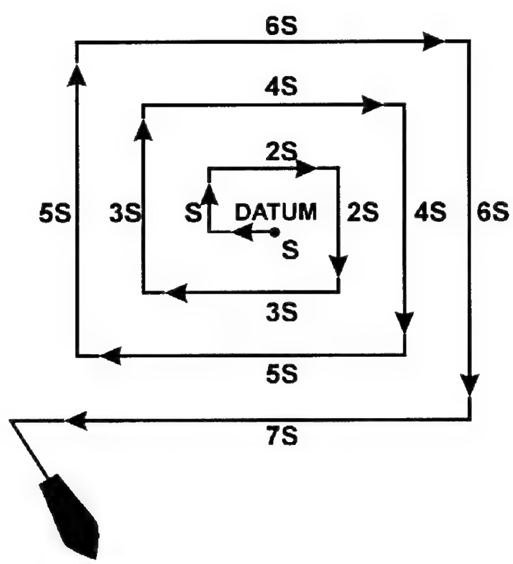

Figure 4-7. Expanding Square Search.

With this type of pattern, the search platform would necessarily pass directly over (or through) the center of the distribution where the probability density was highest. An inverse cube law detector following an infinitely long search leg passing through datum would cut a "groove" or "trough" right in the center of the search object location probability density distribution. Koopman provided a formula for estimating the optimum offset for the next leg in the opposing

direction by computing the optimum offset of a second infinitely long search leg parallel to the first. When Koopman's formula is converted to express this offset, which will become the track spacing, S, in terms of the effective sweep width, W, and the total probable error of position, E, the following is obtained:

$$S \approx \frac{2}{3}\sqrt{W \times E}$$

Once S is known, then the coverage for the first slab may be computed as W/S. Ensuring that the track spacing is optimal for the first few legs of the expanding square ensures that the search plan will be at least nearly optimal. Koopman then assumed that the resulting coverage would be used for all subsequent searches. However, this formula was not used in the CSPM. In fact, very little explicit guidance regarding either coverage or effort allocation was provided.

4.5.2 Koopman vs. Classical Search Planning "Safety Factors"

Although the CSPM provided no explicit guidance, there has been a long-standing, if unwritten, element of U.S. Coast Guard search planning doctrine that strongly favors using a coverage of 1.0 whenever possible. In searches that employ parallel tracks, a coverage of 1.0 implies that the track spacing equals the sweep width ($S = W$). Assuming S and W are equal in equation [4-9], we can compute the values of the first six search radii s_k in terms of the total probable error E to get the following results:

Table 4-1. Koopman Search Radii.

Search Number	Search Radius
1	0.75E
2	1.30E
3	1.68E
4	1.99E
5	2.26E
6	2.50E

The CSPM also contains a method for determining the search radius in terms of E for each search in a series of uniform coverage square searches over an initially circular normal probability density distribution. The values given for use with the CSPM are:

Table 4-2. CSPM Search Radii.

Search Number	Search Radius
1	1.1E
2	1.6E
3	2.0E
4	2.3E
5	2.5E

The coefficients of E listed in Table 4-2 are called "safety factors" in the CSPM. Comparing Tables 4-1 and 4-2, one cannot help but notice that rounding the last three entries in Table 4-1 to two significant digits produces values that are identical to the last three entries in Table 4-2. Furthermore, the second CSPM search radius is still relatively close to the third of Koopman's radii. However, while the first CSPM radius is between the first and second of Koopman's values, there is no obvious relationship between the CSPM first search radius and either of Koopman's first two radii.

4.5.3 Optimal Search Radius for Maximum POS

Upon close examination, it can be seen that Koopman's line of reasoning presents some problems if the objective is to maximize the probability of success. Let us examine the first search represented by the highest and smallest slab in Figure 4-5. To ensure the coverage at datum is optimal, the upper surface of this slab is made tangent to the peak of the paraboloid illustrated by Figure 2-2. It is fair to ask whether applying the amount of effort required to cover this square at a coverage of 1.0 produces the maximum possible POS. To compute the POS, it is necessary to know how much probability is contained in the square. The radius of Koopman's first square is found to be 0.88623 standard deviations. This value may be found either by converting from $0.75269E$ as computed by equation [4-9] to standard deviation or by using equation [3-1] directly. The amount of probability under the normal "bell" distribution between plus and minus 0.88623 standard deviations is 0.62456. The joint probability that both the x-and y-coordinates of the search object's location will be within this range, i.e., the probability that the search object is in the square area, is 0.62456^2 or about 0.39. Assuming an inverse cube law of detection and a pattern of parallel sweeps at coverage 1.0, the POD from equation [4-6] will be about 0.79. Computing the POS as POC × POD produces $0.39 \times 0.79 = 0.3081$ (30.81%).

The amount of search effort required to cover Koopman's first square at coverage 1.0 is simply the square's area or about $(2 \times 0.75269 \times E)^2 = 2.26617E^2$. It is possible to compute the inscribed radius of the optimal search square that yields the maximum POS for this amount of effort. However, the computations cannot be done directly. Numerical techniques must be invoked. The most straightforward solution is to write a computer program that will compute POS values for squares of successively larger (or smaller) size (inscribed radius). This will produce points on a POS vs. radius curve. A method of successive approximations may then be applied to locate the maximum value of this curve. When this is done, it is found that the maximum attainable POS is about 0.31641 (31.641%)—when the radius is about $0.845E$ or, equivalently, 0.9949σ. Since the level of effort was held constant, the coverage of this larger square must necessarily be less than 1.0. The coverage of the larger square would only be about 0.8, yielding a POD of about 0.68397. The probability of containment (POC) for the larger square is about 0.4626. The product of these last two values yields the maximum POS figure given above.

Although the difference in POS values is small (less than one percentage point), this exercise shows that Koopman's allocation of effort was not perfectly optimal in terms of maximizing the POS. However, we seem to be no closer to resolving the mystery surrounding the CSPM's first search "safety factor" of 1.1. Perhaps if we perform a test on the CSPM's recommended first

search square similar to the one we just performed on Koopman's first search square, something of interest will come to light.

4.5.4 Rationales for the CSPM First Search Radius

The CSPM's recommended first search radius is $1.1E$, or about 1.3σ. Using the same method as above, we find that the POC of a square with this inscribed radius is 0.64762. Assuming a coverage of 1.0 with a consequent POD of 0.79 as before, we find that the POS is about 0.51162 (51%). Using the same numerical technique and computer program as before, the optimum radius for $4.84E^2$ units of search effort is found to be $1.09E$, yielding a coverage of about 1.02, a POC of 0.64099, a POD of 0.79889, and a POS of about 0.51208 (51%). It is clear that the CSPM radius is very, very close to optimal for the amount of effort needed to search the recommended square at a coverage of 1.0.

Noting that the POS values produced are just over 50 percent, i.e., since following the CSPM's first search recommendation produces a slightly better than 50-50 chance of locating the search object, it seems reasonable to ask whether this might have been the true objective of the original developers. To obtain a POS of exactly 0.5 using a coverage of 1.0 requires that the square search area contain about $0.5/0.79 = 0.63291$ or about 63 percent of the distribution. The corresponding square has an inscribed radius of 1.26877σ or $1.07759E$. Checking this figure with the computer program produces an optimal search radius for $(2 \times 1.07759E)^2$ units of search effort that is about $1.073E$. The coverage of this slightly smaller square would be 1.00857, the POC would be 0.62980, the POD would be about 0.794, and the POS would be 0.50006. Note that rounding either of the coefficients of E to two significant digits, the CSPM's first search "safety factor" results.

Finally, just for the sake of completeness, we determine the optimal radius for the first uniform coverage 1.0 square search over a circular normal probability density distribution on search object location. That radius is just under $1.065E$, producing a POC of about 0.62445 and a POS of about 0.49331. This is the most efficient coverage 1.0 search in the sense that it produces the highest POS per unit of effort expended. Again, the coefficient of E, when rounded to two significant digits, equals the CSPM's first search "safety factor."

It is not known whether the objective of the original developers of the CSPM was to

1. Find the optimal inscribed radius for the most efficient uniform coverage 1.0 square search,

2. Find the inscribed radius for the most efficient uniform coverage allocation of effort over a square area that would ensure the chances for locating the search object with the first search were at least 50-50, or

3. Find the inscribed radius for a uniform coverage 1.0 square search that would ensure the chances for locating the search object with the first search were at least 50-50.

The latter approach, accompanied by a test to ensure that such an allocation of effort was not seriously sub-optimal, seems to be the most likely approach since it is the easiest to compute. (We must not forget that in 1957 computers were much slower, much more expensive to use, and

much less accessible than they are today. It is very unlikely that the developers of the CSPM had the use of any computers. Not only was the CSPM developed for search planners whose computing resources were limited to paper, pencil and perhaps a slide rule, all the computations done during the ***development*** of the CSPM itself were probably also done by hand.) Within the precision of the CSPM methodology and real-world operations, any of the above approaches would have produced the same result. It seems virtually certain that the CSPM's first search "safety factor" is based on efficiently attaining a desired first search POS of about 50 percent.

Assumption 4: The classical search planning method (CSPM) assumes that the probability of success (*POS*) for the first search should be about 50%, and that this value should be attained with a uniform coverage square search in the most efficient manner possible. These requirements are met very closely by using a "safety factor" of 1.1 to determine the size of the square search area and by covering that area with sufficient effort to attain a coverage of 1.0.

4.5.5 Subsequent CSPM "Safety Factors"

Now that we believe we have discovered the rationale behind the CSPM's first search "safety factor," we can ask whether, given that the first search is performed as recommended but fails to locate the search object, the second search "safety factor" produces an optimal uniform coverage 1.0 square search area. If no searching had been done, the POC for a square having an inscribed radius of $1.6E$ is easily computed to be 0.88424 using the same technique as above. However, if the first search covered a square with an inscribed radius of $1.1E$ at a coverage of 1.0 producing a POS of 0.51162, then the un-normalized POC remaining in the larger square is the difference $0.88424 - 0.51162 = 0.37262$. A uniform coverage 1.0 search would produce a POS for the second search of 0.29437. This is, in fact, an almost perfectly optimal result for $10.24E^2$ units of search effort.

A similar analysis shows that the CSPM's third search "safety factor" of 2.0 is likewise almost perfectly optimal for a coverage of 1.0 ($16E^2$ units of search effort), given the failure of two earlier coverage 1.0 searches of the CSPM's recommended first and second search squares. We can see by consulting Tables 4-1 and 4-2 that at this point, Koopman's original fourth search radius and the CSPM's third search radius are in very close agreement. However, when we look at the CSPM's fourth and fifth search radii, which are also in close agreement with Koopman's fifth and sixth search radii, we find that they do not maximize POS for the levels of effort implied by coverage 1.0 searches.

Table 4-3. Comparison of Recommended Search Radii.

Search Number	Koopman Radius
1	$0.75E$
2	$1.30E$
3	$1.68E$
4	$1.99E$
5	$2.26E$
6	$2.50E$

Search Number	CSPM Radius	Optimal Radius
-	-	-
1	$1.1E$	$1.09E$
2	$1.6E$	$1.59E$
3	$2.0E$	$2.02E$
4	$2.3E$	$2.40E$
5	$2.5E$	$2.74E$

In light of the preceding discussions, the following interpretation of Table 4-3 seems as likely as any: The first three CSPM "safety factors" were chosen to produce, for an inverse cube detection function, very nearly optimal (maximum POS) search areas for a sequence of uniform coverage 1.0 concentric search squares over a circular normal search object location probability density distribution. The third CSPM radius agreed with the fourth Koopman radius. The latter was based on uniform coverage 1.0 search squares that kept the "cumulative (random search) coverage" at the datum position optimal while producing a "piled slab solid" whose volume (search effort) equaled that of the optimal paraboloid. Rather than attempt to find the optimum (maximum POS) inscribed radius for additional uniform coverage 1.0 square searches, a task that was becoming increasingly difficult, especially without a computer, the decision was made to revert to Koopman's radii for the fourth and fifth "safety factors." Thus the two techniques were "blended" together at the point where both produced essentially the same result. The aesthetically pleasing sequence of differences between successive CSPM "safety factors" that resulted (0.5, 0.4, 0.3, 0.2) may also have influenced the original CSPM developers.

Assumption 5: The classical search planning method (CSPM) assumes that for all searches, the track spacing will equal the sweep width, i.e., the coverage will be 1.0.

Assumption 6: The classical search planning method (CSPM) assumes that for all searches, the size of the recommended search square will be governed by the prescribed "safety factors," and the amount of search effort available and expended will equal the amount required to cover the recommended search square at a coverage of 1.0.

4.6 CLASSICAL SEARCH PLANNING IN PERSPECTIVE

At the time of its development in 1957, the classical search planning method (CSPM) was an outstanding example of a brilliant and pragmatic application of search theory to the practical problem of searching for distressed persons and craft at sea. Despite the true complexity of the marine SAR search problem, the method was simple, relatively quick and easy to use, and generally well matched to the available technologies of the day.

Although it served the marine SAR community well for many years, the CSPM was far from perfect. Each of the six major assumptions on which the CSPM is based requires a number of underlying assumptions, many of which are often a poor match for the SAR situation at hand. In addition, technology has changed dramatically since 1957 in many areas that affect SAR and search planning. We will examine the assumptions and technology changes briefly in the paragraphs that follow so we can identify some of the CSPM's basic weaknesses. A few of these weaknesses could have been addressed, at least partially, in the original development without adding significantly to the complexity or computational burden of the CSPM, i.e., a few were preventable. Most could not be addressed without the aid of computers and other technological developments as well as more recent developments in search theory.

4.6.1 Pre-Search Probability Density Distributions

There are a number of "simplicity constraints" imposed on the CSPM. These are needed to keep the computational burden on the search planner within reasonable limits. However, each of these presents problems when it comes to representing actual SAR scenarios.

4.6.1.1 Incident Position

If a specific single distress position (and time) is given or one can be estimated, along with its probable error, a circular normal probability density distribution of possible distress positions is often a reasonable assumption. Figure 4-8 below shows the corresponding probability map.

0.03%	0.38%	0.88%	0.38%	0.03%
0.38%	4.97%	11.59%	4.97%	0.38%
0.88%	11.59%	27.05%	11.59%	0.88%
0.38%	4.97%	11.59%	4.97%	0.38%
0.03%	0.38%	0.88%	0.38%	0.03%

Figure 4-8. Circular Normal Probability Map.

However, there are many SAR situations where there is insufficient information to estimate either a specific incident position or a specific time of the incident.

A classic situation where a single distress position and time cannot be reasonably estimated is when a vessel or aircraft becomes overdue or unreported. One example is when a vessel or aircraft in transit between two points fails to arrive at its destination (overdue) or fails to report as scheduled during the transit (unreported). Another example is when a vessel, such as a fishing vessel, fails to report as scheduled during operations in some general area (e.g., the Grand Banks) or fails to return to port. A slightly different situation arises when a vessel or aircraft issues a distress call but either does not provide a position or the receiver is unable to copy the position. In this case, the time of the incident is known and sometimes a position can be estimated if the identity and intended route of the distressed craft can be ascertained. However, distress situations often arise in circumstances where the distressed craft was unable to maintain its intended route and/or schedule, so positions estimated in this fashion must be treated with extreme caution, as they can be very misleading.

Although some attempts have been made to extend the CSPM so it would cover situations such as those described above, they are very difficult to implement manually and often cause more problems than they solve. In an overdue transit scenario, the initial search object location probability density distribution clearly cannot be represented by a single circular normal distribution about a specific point at a specific time. Instead, the distribution is spread along the missing craft's intended track and also spread over some period of time. One extension of the CSPM calls for choosing a number of possible incident positions and times along the intended track, computing drift updated datums and probable errors for each, and enclosing the resulting circles and the areas between them with some arrangement of rectangular search areas. In a manual method, only a few datums may be computed in this fashion and the results may not represent the situation at the commence search time very well. It is virtually certain that any manually developed search plan for such situations will be substantially sub-optimal as compared to plans that could be developed with appropriate computer assistance.

4.6.1.2 Mean Drift Estimation

The CSPM assumes a homogeneous environment in space and time where variations from the mean drift vector are random and have a circular normal probability density distribution. This in turn requires that all the components of the drift vector (leeway, wind current, sea current, etc.) all share these same properties.

The assumption of a homogeneous environment is most realistic in environmentally stable areas such as where the trade winds blow over the open ocean far offshore. In most other parts of the world, wind and current often vary significantly and systematically (vice randomly) over space and time on scales that are small compared to the areas that are covered by initial distributions and the lengths of the drift intervals. For example, a probable incident position error of only five nautical miles implies that half the distribution lies within five nautical miles of the incident position while all but a tiny fraction of the other half lies between 5 and 15 nautical miles of the incident position. This means the distribution covers more than 700 square nautical miles. Furthermore, the size of the distribution will grow larger with time as drift and drift uncertainties are taken into account. Scenarios involving overdue or unreported craft typically cover much larger areas and may involve significant differences between the earliest and latest possible incident times. In addition, drift intervals are often 24 hours in length or longer. Significant changes in wind speed and direction often occur over much shorter spans of time. Near shore, tidal currents are known to vary substantially and systematically over short distances and time spans.

For the CSPM, the assumption of a homogeneous environment was an inescapable approximation. Only one drift trajectory was being computed. If this single trajectory was to be representative of all possible trajectories, i.e., the mean of all possible trajectories, the environment had to be stable and homogeneous. In 1957, such an assumption was not an unreasonable approximation offshore due to the lack of detailed environmental data. There were no supercomputers running sophisticated models of the atmosphere and ocean, no satellites gathering and relaying back to earth masses of environmental data every day (in fact, there were no man-made satellites at all), no offshore special-purpose weather buoys (just synoptic observations from a few light ships, ocean station vessels, and vessels of opportunity), no weather radar, etc. Environmental data for search planning was restricted to small-scale paper

atlases, pilot charts, and possibly a few observations at or near the scene from craft of opportunity. Search planners really had no choice but to assume the sparse, low-resolution data available to them was homogeneous over large areas and spans of time. As a result, there was a reasonably good match between the CSPM and the quality of the available environmental data.

Today, a number of organizations routinely provide high-resolution environmental data that resolve many significant environmental features that were previously invisible to search planners.

When currently available environmental data products are displayed, any observer familiar with the earlier data sources and the CSPM can quickly discern some important differences between the situation today and that in 1957:

Just one or a handful of computed drift trajectories cannot realistically represent the possible drift trajectories and the consequent probability density distribution of possible search object positions that must result from the detailed environmental data now available, even in cases where the incident position and time are accurately known. Furthermore, such distributions are very, very unlikely to be circular normal or have any other simple form, even if the incident position probability density distribution was circular normal at the time of the incident.

When high-resolution data are available, computed datum positions often become very sensitive to the exact placement of the estimated incident position. Even a small error in this estimate could make a difference of several miles in the computed datum position over intervals of less than 24 hours.

Even if one or a handful of computed drift trajectories could realistically represent the distribution of possible drift trajectories, the data are so finely gridded in some places that drift sub-intervals as short as one hour would be required to use the available data effectively and avoid skipping over data points without using them. This would be a 24-fold increase in the computational burden placed on the search planner.

Figure 4-9 below illustrates how a typical CASP probability map might look after a day or so of drift.

In short, the CSPM, and its extensions (to be discussed later), are very poorly matched to the quality and resolution of the environmental data now available. In fact, the match is so poor as to render the CSPM and its extensions virtually useless without computerization, and even then there are other, far more effective, ways to use the same computing resources to aid the search planning process.

0.16%	0.33%	3.25%	0.97%	0.42%
0.51%	1.45%	7.44%	5.37%	3.33%
2.43%	3.92%	7.77%	7.50%	5.15%
3.48%	4.43%	3.27%	6.29%	5.49%
2.12%	6.07%	3.33%	2.00%	1.52%
1.40%	3.88%	2.92%	2.05%	1.75%

Figure 4-9. An Example of a CASP Probability Map.

4.6.1.3 Drift Error Estimation

Perhaps the weakest link in the CSPM right from the beginning has been the way in which drift error has been estimated. This weakness has now become critical due to improvements in navigation technology. The Global Positioning System (GPS) now allows virtually all persons and craft to determine their positions to within a few meters with a quite modest investment in a GPS receiver. Recall that the total probable error of the search object's position is a function of the incident position error, the search craft's position error and the drift error. The reasonable cost and widespread use of GPS has greatly decreased the number of cases where the uncertainty in either the distressed craft's reported position or the search craft's position is significantly different from zero. In an increasing number of cases, the drift error is the only significant source of error in the search object's location and the only significant factor affecting optimal search area size. Therefore, obtaining accurate and statistically valid estimates of drift error has become imperative for planning searches. There are several aspects of the CSPM drift error estimate that need closer examination if we are to understand the nature of this problem.

Recall that the CSPM began by estimating the probable drift error as one-eighth of the distance drifted. Roughly speaking, this implies that for half of all drift distances computed by search planners, the computed distance drifted would be within 12.5 percent of the correct value, regardless of the amount of time the object has been adrift. When one considers the sparseness and low quality of the data the search planners were using, this was a very optimistic assumption. Even today, it would be difficult to find a professional oceanographer with a long history of studying a particular region of ocean who would claim to be able to accomplish such a feat even in the area under study. The problem, stated in terms of physics, is to predict the motion of a

tiny solid object suspended at the turbulent interface between two huge masses of dissimilar fluids (the ocean and the atmosphere). This is a very difficult thing to do with any accuracy.

If there were four equal sources of error, each expressed as a fraction of its respective mean value, then Theorem 3 of paragraph 4.4.3 tells us that each of these would need to have a probable error equal to one-sixteenth (1/16) of its mean to produce a total probable error equal to one-eighth (1/8) of the resultant mean vector's magnitude. Environmental data that is within 6.25 percent of the correct value half the time is very accurate indeed. Even today it would be hard to find data of this quality that would be available to search planners. In short, the 1/8 factor for drift error was unrealistically small. This factor was later increased to three-tenths (3/10, 0.3), but even this may have been optimistic. For four equal sources of error, each would have to be within 15 percent of its mean value half the time to produce a resultant probable error of 0.3.

In addition to the assumed value of the error, or "confidence" factor, the technique itself is a very poor way to estimate probable error. Estimating probable error as a fraction of the mean value forces the error estimate to be correlated with the mean value. A small mean produces a small error estimate while a large mean produces a large error estimate. Very often the size of the error is largely independent of the mean value. For example, in the trade winds, the mean wind might be ENE/15 knots with a probable error of only one or two knots for several days. On the other hand, a fast-moving gale passing over the Grand Banks might cause the wind to "box the compass" in a 24-hour period, causing the vector average wind over the period to be close to zero. However, the probable error of that estimate could easily be 15 or 20 knots or more. When dealing with even just a single environmental parameter, estimating the probable error as a fixed fraction of the mean value is, at best, very unreliable.

Drift estimates typically depend on several environmental parameters (e.g., winds and currents) and on our knowledge of how other elements of the problem respond to them (e.g., wind current, search object leeway). The direction and rate of drift are estimated by adding leeway and current vectors. Since the wind and current are often largely independent of one another, any combination of leeway and current are possible. If the leeway is in the same direction as the total water current (TWC), then the drift velocity, drift distance and CSPM-estimated drift error may be large. If the leeway and TWC are opposed, then the drift velocity, drift distance and CSPM-estimated error are likely to be small. The actual amount of error in the drift estimate should depend on the quality of the environmental data and the quality of our knowledge of how search objects and the ocean's surface react to environmental forces. Certainly the relative directions of the winds and currents should not affect the quality of the drift estimate.

Figures 4-10a to 4-10f show the results of simulations where the probable errors of the environmental data and the methods for computing the search object's response to them were held constant, the mean TWC was held constant, the mean magnitude of the leeway was held constant, but the leeway direction "boxes the compass" in 60 degrees increments in a clockwise direction as one progresses through the figures. The initial position error at IP was taken to be 0.1 NM. Five hundred points were distributed around the IP at random according to a circular normal probability density distribution. A mean sea current was entered with a probable error of 0.3 knots. The mean leeway and mean wind current were computed from the mean wind that

was input. Each of these was assumed to have a probable error of 0.3 knots. Using Theorem 3 from section 4.2.3, the total probable drift velocity error was computed to be 0.5196 knots, resulting in a probable drift position error after 24 hours of 12.47 NM (Theorem 4). To demonstrate the validity of this value and compare it with the probable error computed by the CSPM method, the following Monte Carlo technique was used: For each of the points initially distributed around the IP, a sea current error vector was chosen at random and added to the mean sea current to get a sample sea current. Sample leeway and wind current vectors were obtained in the same way. All samples were independently drawn. These sample vectors were then added to get a sample drift velocity direction and speed in knots. The sample drift velocity was multiplied by the length of the drift interval (24 hours in this case) to get a sample drift direction and distance whose position was then plotted. In this fashion, the results of 500 independent drift trajectories were computed.

In each figure below, the heavy black circle shows the probable drift error computed from the probable errors of the inputs to the drift computation. The red circle shows the CSPM estimate of the probable error using a "confidence factor" of 0.3. The probability of containment (POC) values given in the tables under the figures are based on counting the number of points contained in the circles shown.

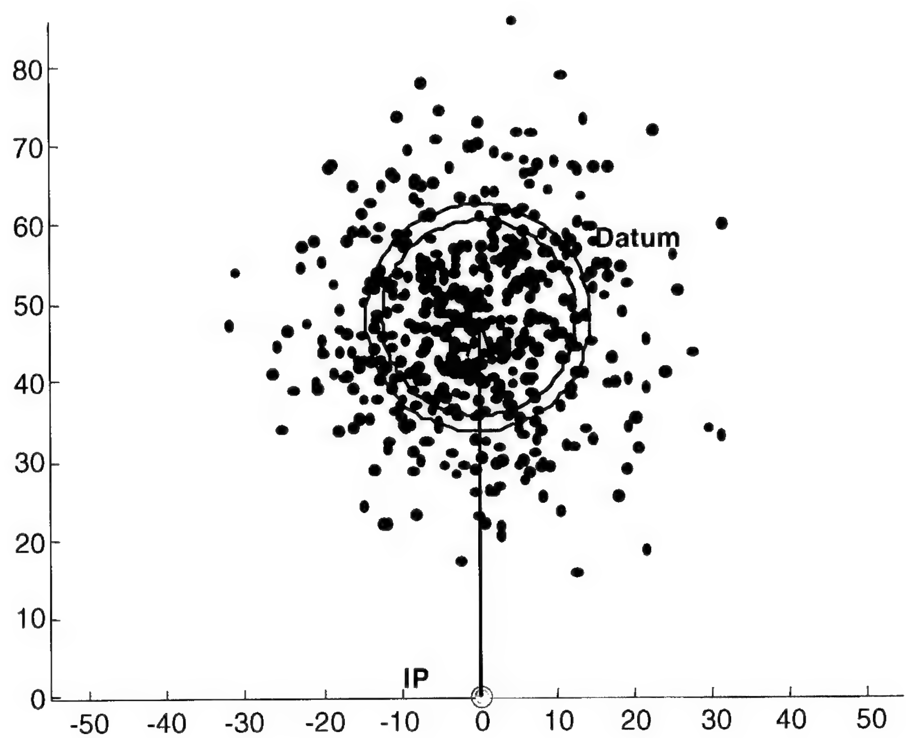

Figure 4-10a. Monte Carlo Distribution.

Average values: Wind 180T/20 kts, TWC 000T/1.0 kt, Leeway 000T/1.0 kt

	E	Probability Contained in Circle
Present CSPM Method	14.40 NM	58.2%
Proposed Method	12.47 NM	50.0%

In Figure 4-10b below, both methods happen to produce identical results.

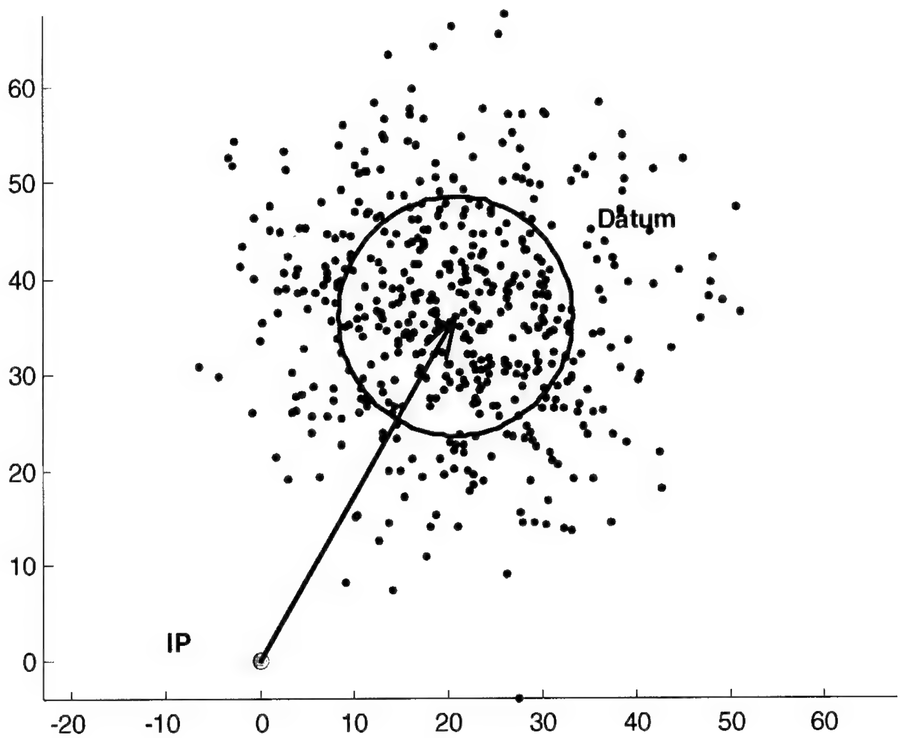

Figure 4-10b. Monte Carlo Distribution.

Average values: Wind 240T/20 kts, TWC 000T/1.0 kt, Leeway 060T/1.0 kt

	E	Probability Contained in Circle
Present CSPM Method	12.47 NM	49.0%
Proposed Method	12.47 NM	49.0%

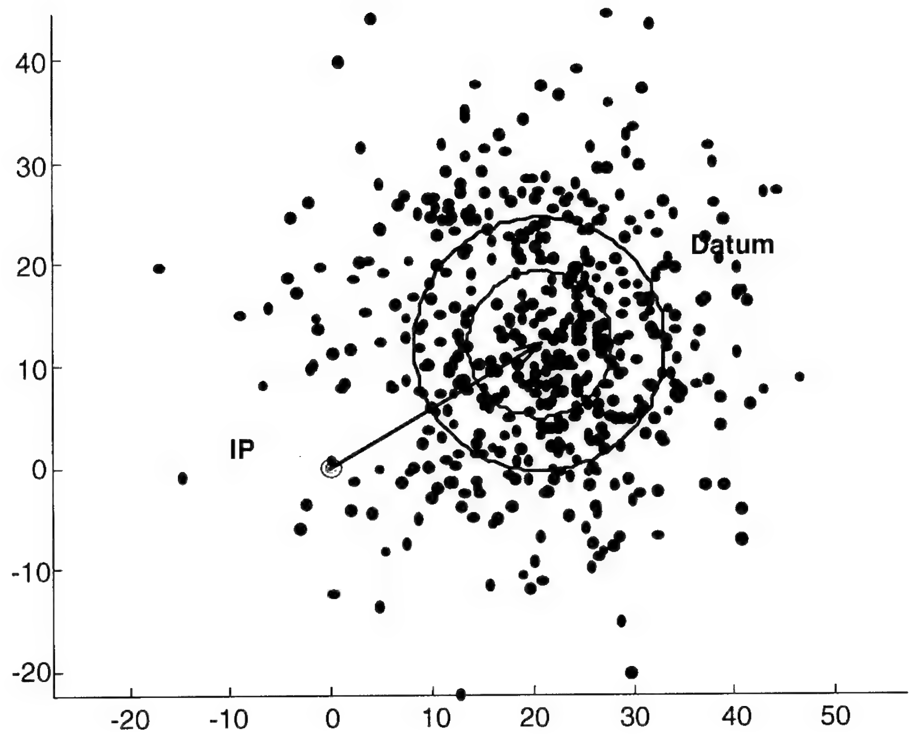

Figure 4-10c. Monte Carlo Distribution.

Average values: Wind 300T/20 kts, TWC 000T/1.0 kt, Leeway 120T/1.0 kt

	E	Probability Contained in Circle
Present CSPM Method	7.20 NM	20.6%
Proposed Method	12.47 NM	49.0%

Note that in Figure 4-10d, the probable error of position using the present method is ***not*** missing. Since the average leeway and average total water current have exactly the same magnitudes but opposite directions, the net mean distance drifted is zero. This means Equation [3-7] will compute a total probable drift error of zero. The average initial position and the datum position are in the same place in Figure 4-10d.

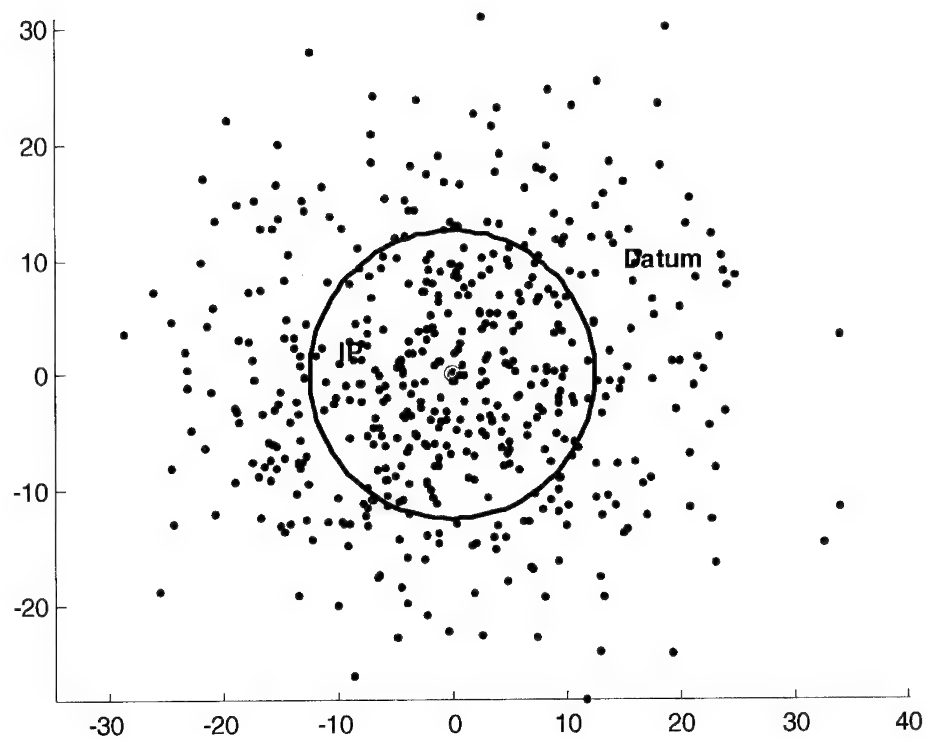

Figure 4-10d. Monte Carlo Distribution.

Average values: Wind 000T/20 kts, TWC 000T/1.0 kt, Leeway 180T/1.0 kt

	E	Probability Contained in Circle
Present CSPM Method	0.1 NM	0.2%
Proposed Method	12.47 NM	48.0%

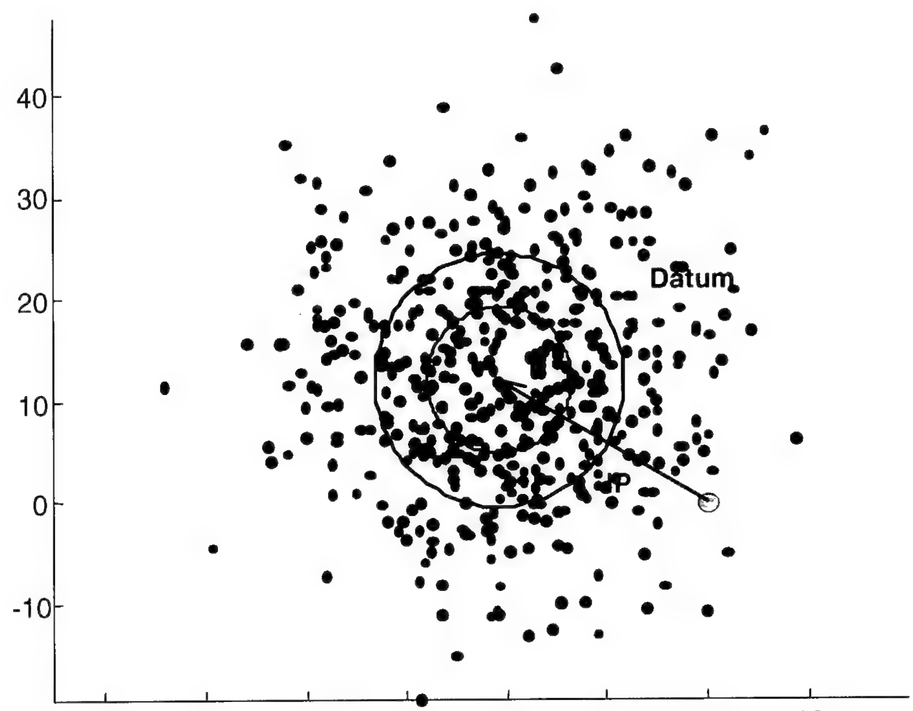

Figure 4-10e. Monte Carlo Distribution.

Average values: Wind 060T/20 kts, TWC 000T/1.0 kt, Leeway 240T/1.0 kt

	E	Probability Contained in Circle
Present CSPM Method	7.2 NM	20.8%
Proposed Method	12.47 NM	50.0%

In Figure 4-10f below, both methods again happen to produce identical results.

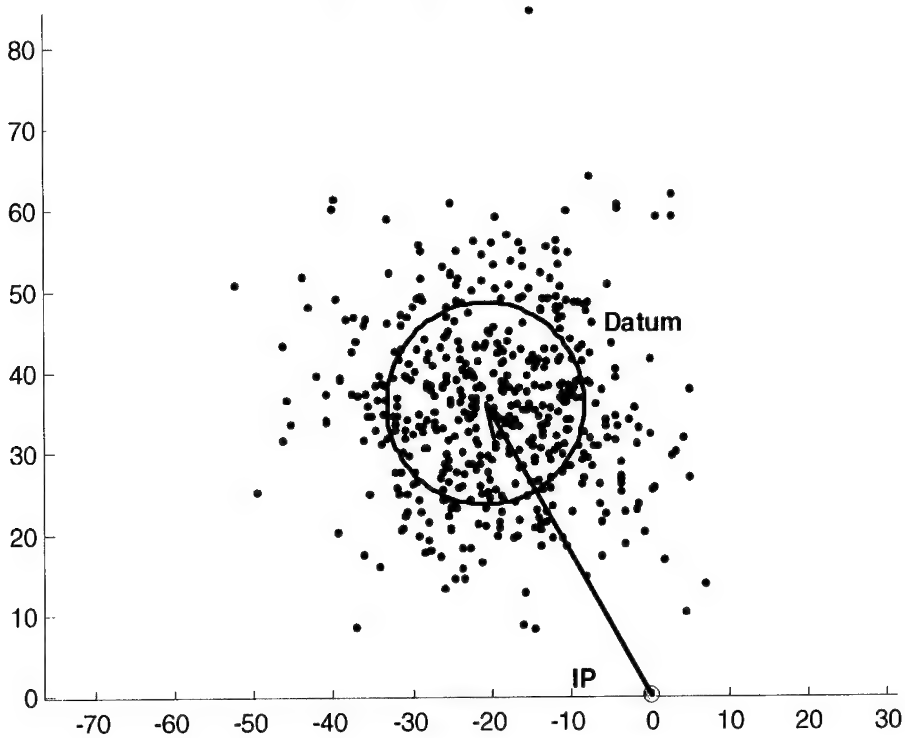

Figure 4-10f. Monte Carlo Distribution.

Average values: Wind 120T/20 kts, TWC 000T/1.0 kt, Leeway 300T/1.0 kt

	E	Probability Contained in Circle
Present CSPM Method	12.47 NM	49.6%
Proposed Method	12.47 NM	49.6%

Note that the CSPM estimate of the total probable error varies from virtually zero to a value somewhat greater than the total probable error that was computed from the actual "component" probable error values on the parameters used in the simulation. The validity of the latter method is confirmed by the fact that the number of points contained in the black circles is always very close to 50 percent of the total number of points in the simulation. This is a graphic example of why estimating the probable error of the sum of several vectors as a fixed fraction of the resultant's magnitude is extremely unreliable.

Finally, as the accuracy of environmental data and models of search object drift motion improve, the probable errors in computed drift values will decrease. Such improvements should result in an appropriate decrease in the probable drift error that in turn will lead to a decrease in the total probable error of position. Ultimately, this means the size of the optimal search area for a given level of effort will decrease. More concentrated searching in smaller areas that have a high probability of containing the search object will lead to earlier search object detection and more lives saved with fewer resource hours. In short, such improvements will take some of the "search" out of search and rescue. However, there is no way to realize these benefits as long as the drift error is estimated as a fixed fraction of the distance drifted.

Historical note: At some point prior to 1973, the method of computing drift error was modified. It was recognized that a drift trajectory could turn back toward the incident position at some point, thus shortening the net distance drifted since the object started drifting. Taking one-eighth of this distance would have caused a reduction in the size of the assumed circular normal probability density distribution when common sense and the laws of physics indicated it should continue to expand. The "patch" to fix this problem was to compute probable drift error as one-eighth of the distance drifted since the time of the last computed datum plus the sum of all previously computed drift errors. This was not statistically correct, but it did have the effect of forcing the drift error to increase monotonically.

4.6.1.4 Summary of Pre-Search Distribution Issues

A single incident position with a circular normal probability density distribution of position errors is inadequate for representing the initial distributions for many SAR situations. This portion of the CSPM cannot be extended to include other types of distributions in a manner that is simple to implement, statistically correct, and consistent with the remaining portions of the CSPM.

In most of the world, a single mean drift trajectory with a circular normal probability density distribution of drift errors is inadequate for representing the distribution of possible drifting object positions after more than a few hours adrift, even when the initial distribution of possible incident positions has a circular normal probability density. Often the post-drift datum position is very sensitive to small variations in the incident position estimate.

The CSPM technique for estimating probable drift error is fundamentally flawed, making it inadequate, inappropriate, and unreliable. This has become a critical problem in recent years as modern navigation technology, especially GPS, has virtually eliminated both distressed craft and search craft contributions to the total probable error of position. This has left drift error as the sole significant contributor to search object positional uncertainty, a primary determinant of search area size. The CSPM drift error estimation technique also does not provide a mechanism to account for improvements in either the quality of environmental data or the accuracy of drift motion models by reducing the drift error estimate appropriately.

The CSPM's method for establishing initial search object location probability density distributions is poorly matched to reality and poorly matched to current technology. Current technology makes much more realistic estimates of initial distributions possible.

4.6.2 The Detection Function

In 1957 the developers of the CSPM chose to use the detection function derived from the inverse cube model of visual detection developed by Koopman [1946, 1980]. This choice carried with it a large number of implied assumptions, some of which are clearly invalid for SAR. Koopman had a ready answer for this problem. Making reference to a figure like Figure 4-11 below, he wrote,

> "At one extreme is the case of the definite range law, at the other the case of random search. All actual situations can be regarded as leading to intermediate

curves, i.e., lying in the shaded region. The inverse cube law is close to a middle case, a circumstance which indicates its frequent empirical use, even in cases where the special assumptions upon which its derivation was based are largely rejected."

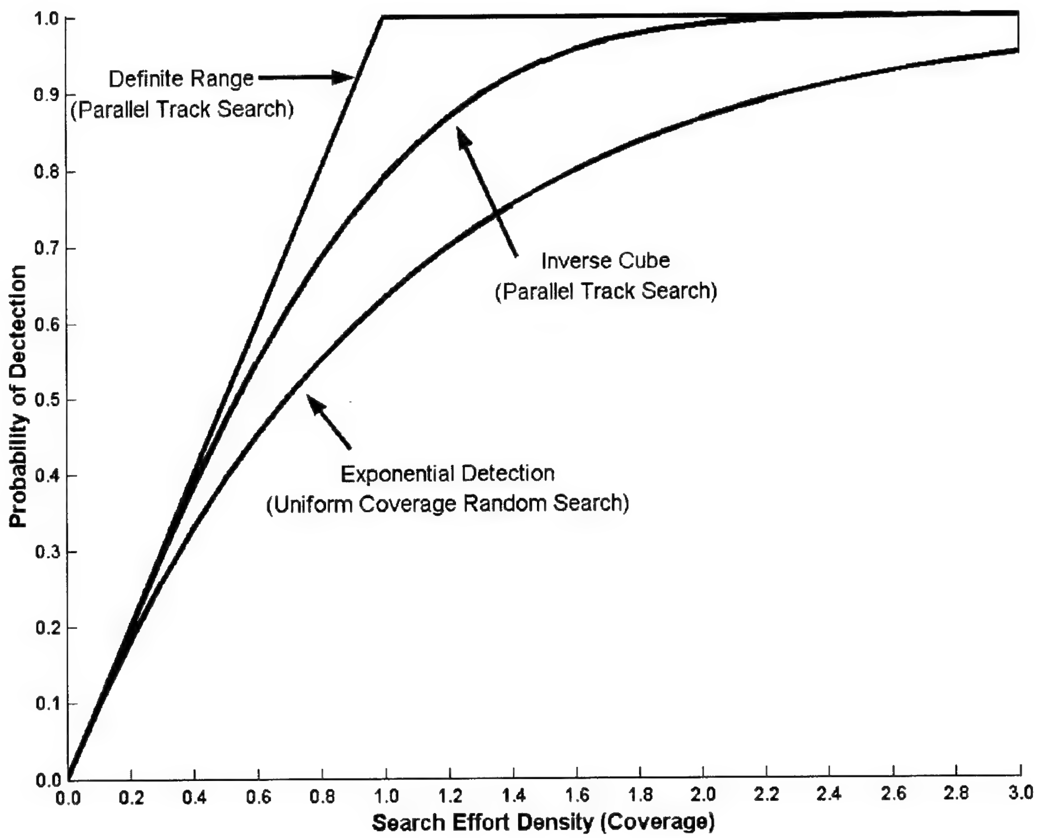

Figure 4-11. Three Detection Functions.

Nevertheless, we shall examine these special assumptions with a view toward determining whether they may be safely ignored in the context of SAR. There are two basic categories of assumptions to examine. First, there are the assumptions directly related to the nature of visual detection. These are the assumptions that led to the inverse cube model of instantaneous detection by the unaided human eye. Second, there are the assumptions related to the motion characteristics of both search platforms and search objects.

4.6.2.1 Inverse Cube Model of Instantaneous Visual Detection

The basic assumptions of the inverse cube model of instantaneous visual detection were stated in paragraph 4.3.1 above. We will examine each and comment on its applicability to SAR.

The first assumption, that the observer is at some height *h* above the ocean's surface, seems reasonable since virtually any searcher will have some "height of eye" above the surface.

However, there are two problems. First, the assumption makes sense in the context of Koopman's geometric rendition of the model (Figure 4-1) only for airborne search platforms. Although most SAR searches involving large areas are conducted from aircraft, many localized searches are conducted from vessels where other factors, such as horizon distance, sea state, etc., interfere with Koopman's assumed geometry. Second, Koopman makes the approximation that

$$\gamma = \frac{kh}{s^3} = \frac{kh}{\left(h^2 + r^2\right)^{3/2}} \approx \frac{kh}{r^3}.$$

He justified this approximation by claiming that the detection range, r, is much larger than h in the majority of cases. His claim was probably reasonable for the wartime situations he envisioned. However, for small SAR search objects, h is a significant fraction of r, in the majority of cases for airborne search platforms. Koopman's approximation also implies that detection of all search objects lying on the search craft's track is guaranteed (POD = 100%) regardless of the search platform or search conditions. There is ample evidence from U.S. Coast Guard operational SAR experience that this is not the case. Nevertheless, this characteristic of Koopman's approximation contributes significantly to the form and POD values of the corresponding POD vs. Coverage curve for parallel track search patterns. A non-zero value for h will produce a lateral range curve that has a different shape and a lower maximum detection probability than that produced by Koopman's approximation. These changes will tend to move the POD vs. Coverage curve downward toward the "random" search curve.

The second assumption, that the search object is a vessel underway, is clearly false in the vast majority of SAR cases. Vessels of the type Koopman had in mind—namely warships at least the size of a surfaced WWII diesel-electric submarine—are rarely the object of a SAR search. SAR searches usually seek persons in the water, life rafts, or small boats that are adrift.

The third assumption, that the search object is initially detected by sighting its wake, is also clearly false in virtually all SAR cases. Drifting objects, even those with significant leeway, leave no visible wake. The object itself must be sighted for a detection to take place.

The last assumption, that the instantaneous (one-glimpse) probability of detecting a cruising vessel is proportional to the solid angle subtended at the observer's eye by the vessel's wake is likewise clearly false in a SAR context because there is no visible wake present to detect. However, one could make the argument that Koopman's geometry could be applied to the search object itself rather than its wake. Such a modification would at least keep Koopman's model within the realm of plausibility for SAR. Unfortunately, in all the years since Koopman first postulated the inverse cube model of visual detection, it does not appear that anyone has attempted to test his hypothesis either in the laboratory or in field trials.

4.6.2.2 Search Platform and Search Object Motions

Koopman's inverse cube model of instantaneous visual detection standing alone does not provide a way to relate the probability of detecting an object known to be in some area with the amount

of searching effort expended in that area. Certain other assumptions are necessary. These include the following:

1. Either the search object is stationary and the search platform approaches the search object from a long distance away along a straight path at constant speed, passes the object at some lateral range, and continues along the same straight path at the same speed for a long distance beyond the search object, or both the search platform and search object are moving along long, straight tracks at constant speeds for a long time before and after their closest approach, producing the same effect *relative to the moving search object* as in the case of the stationary search object.

2. Areas are searched by means of long, straight, equally spaced, parallel tracks *relative to the search object.*

The first assumption allowed Koopman to develop a *lateral range curve* that expressed the probability an object would be detected by an inverse cube "sensor" as a function of the lateral range or distance between the search platform and the search object at the closest point of approach. Figure 4-2 depicts the lateral range curve for Koopman's simplified inverse cube model.

Koopman's assumption about the basic nature of the relative motion between the search platform and the search object may have been a reasonable one when using aircraft to patrol for enemy ships. Although air navigation was poor by today's standards, the sweep widths were large in comparison to the expected magnitude of the variations from the intended track due to random navigational error. It was also reasonable to expect enemy ships to follow approximately straight tracks at constant speeds.

SAR searches tend to be quite different in character. The search objects are small with correspondingly small sweep widths. Until recently, aircraft navigation errors were often large in comparison to sweep widths. Diverting from the intended search track to investigate sightings was probably more frequent than it was in wartime patrols, adding to the problems of navigation. Finally, drifting objects are subject to the vagaries of the environment and do not tend to move in long straight lines like ships in transit. In recent years a great many satellite-tracked buoys have been deployed for various oceanographic studies. The tracks of these buoys are jagged and often contain "inertial loops" where the buoy exhibits a circular or cycloidal motion as it follows the water in which it is floating. Other phenomena such as rotary tidal currents, warm-core and cold-core eddies spun off from the Gulf Stream as well as its meandering, and other natural phenomena tend to complicate drift trajectories even more. For objects with significant leeway, wind shifts often occur over relatively small time scales, complicating both drift trajectories and the navigational problems of searching aircraft. Nevertheless, the first assumption above is still a good approximation for any single track due to the short time during which the search object can be within the search craft's detection envelope. For example, if the maximum (practical) detection range is 4.5 NM and the search aircraft is moving at 180 knots, then the search object will be detectable for only about three minutes at most, give or take a few seconds to account for search object drift rate. On the other hand, it could take an aircraft several hours to complete its assigned search pattern and this presents a potential problem for the second assumption above.

Most searches over water use some type of parallel track search pattern relative to the earth. Typically, a single aircraft tries to perform a series of parallel sweeps in an assigned search area. Generally it takes several hours to complete a search pattern. With no landmarks for ready reference, and without the most modern of navigational aids, flying a perfect search pattern over the ocean is a very difficult proposition. Figure 4-12 below shows how a search craft's actual track might compare to its intended track. (Very similar examples, and worse, were actually recorded by accurate tracking devices during early sweep width experiments.) If the search object is stationary, then a perfect geographic pattern will remain intact relative to the search object. However, if the object is moving, such patterns are susceptible to both additional random and systematic distortions when viewed from the search object's perspective.

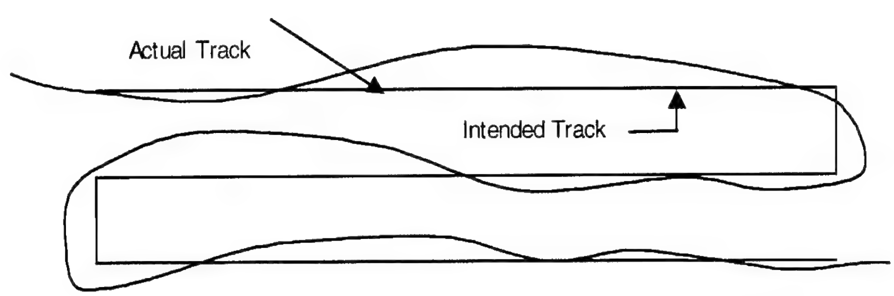

Figure 4-12. Actual versus Intended Search Craft Track.

Although drifting search objects move very slowly when compared to the speeds of searching aircraft, it is not unusual for them to drift several miles while the aircraft are on scene searching. It is also not unusual for an hour or more to elapse between the time when the search aircraft is at a certain point on one leg and the time at which it returns to a similar position on an adjacent leg. Since search objects are moving, often with a significant degree of randomness over such time scales, and often at a significant rate of speed compared to the search craft's *creep rate* perpendicular to its search legs, there is no guarantee that the search object will remain in substantially the same location relative to the search pattern as the search progresses.

Figure 4-13 below shows an example of a systematic distortion of an otherwise perfect search pattern. In this case, the search craft was "creeping" in the same direction as a search object moving with constant direction and speed. The rate of creep was four times the drift rate. To give a concrete example, if the search area was 60 NM long in the direction of the search legs, the track spacing was 4.0 NM, and the aircraft's search speed was 120 knots true speed over the ground, then the creep rate is 8.0 knots. If the search object were drifting in the same direction as the creep at 2.0 knots, then we would have the situation depicted in Figure 4-13.

Note that the search legs are no longer parallel and that about 25 percent of the search area (the portion above the dashed line) is not covered. Any object in the northern 25 percent of the planned search area would have drifted out of the area by the time the search pattern was completed. The average coverage of the other 75 percent is about 1.33 times the intended coverage, but this coverage is less uniform than it would have been if the search legs had remained parallel. The result is a somewhat lower POD for the reduced area than one would expect from a perfect pattern at the increased coverage. If the planned search had been

optimized for maximum POS, the distortion shown would likely result in a substantially sub-optimal POS value.

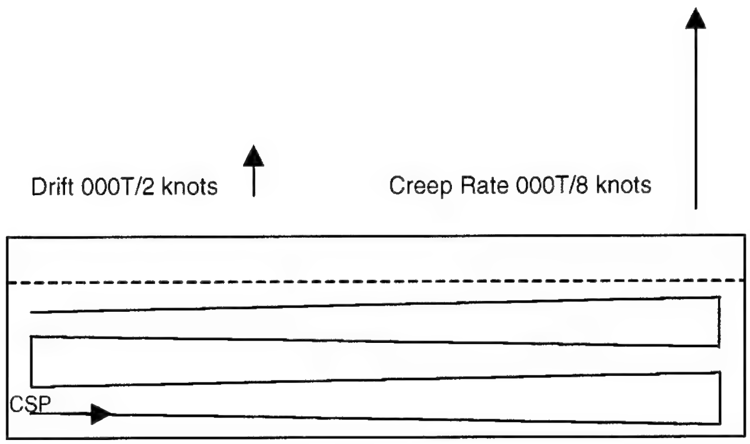

Figure 4-13. Systematic Relative Motion Distortion.

Koopman [1946] developed "barrier" search patterns to address this type of distortion and produce undistorted patterns relative to a moving search object. These were used primarily for defense and blockades where the objective was to prevent the undetected movement of enemy shipping into or out of an area when the speed and direction of such movements were either known or could be predicted with reasonable accuracy. The barrier search pattern is not particularly useful for SAR since in many cases there can be considerable variation in search object drift from one possible location to another within an area that makes, from an operational perspective, a good search area otherwise. Barrier patterns are quite sensitive to differences between the search object's predicted speed, which the barrier is designed for, and the actual search object speed.

Orienting the search legs parallel to the mean predicted direction of drift makes the search pattern much less sensitive to variations in search object speed. The search object's motion may then be addressed in one of two ways. The length of the area can be extended in the down-drift direction far enough to ensure any object in the originally desired search area will not be able to leave the planned search area before the search has been completed. Alternatively, the desired search area can be skewed into the shape of a parallelogram to compensate for the object's motion. Although this type of pattern probably has very limited non-SAR military value, orienting the search legs in the direction of the mean expected drift vector has proven very effective for SAR searches.

Nevertheless, search object motion during the performance of a search remains an uncertain quantity. It is impossible to compensate for all possible drift trajectories in any practical search operation. Therefore, when one is faced with estimating the probability of detecting a search object given that it was in the search area when the search began, but whose exact motion during the search is unknown, one is faced with something approaching "random" searching even when there is compensation for the mean motion. It is almost certain that the POD estimates for past searches have been overly optimistic because the assumption of "perfect" search patterns relative to the search object was not met.

In many past cases, and probably some more recent ones as well, failure to compensate for mean expected drift must have resulted in significant systematic pattern distortions in addition to the non-systematic distortions caused by the random component of the drift motion and search craft navigational error. The latter can also have both random and systematic (when navigational conditions are poor) components. The result would have been grossly inflated assessments of search effectiveness, again because the assumption of "perfect" patterns relative to the search object was false.

4.6.2.3 Summary of Inverse Cube Issues

The primary justifications for using Koopman's inverse cube model of visual detection seem to be the following:

The inverse cube model of visual detection has a plausible mathematical basis.

"Common sense" seems to dictate that an organized parallel sweep search effort with imperfect "real world" sensing should produce results falling between those of "perfect" definite range sensing in combination with perfect parallel track search patterns and those of a uniform coverage "random" search.

The inverse cube POD vs. Coverage detection function is regular and falls roughly halfway between that of the definite range model and that of uniform coverage "random" search.

Under ideal search conditions, U.S. Coast Guard Research and Development Center studies often produced results that were roughly consistent with Koopman's model.

However, it seems that a stronger case can be made for using the so-called "random" search detection (vs. coverage) function for SAR searches in the marine environment.

There is no experimental evidence to support (or refute) Koopman's hypothetical inverse cube model of visual detection, even for wakes of large vessels.

All but the first of Koopman's assumptions leading to the inverse cube model are clearly false for SAR situations. Even that first assumption (that the observer is in an aircraft flying over the ocean at a constant altitude) is false when the search platform is a vessel.

Koopman's simplification of the inverse cube model's lateral range curve produces an unrealistically high POD at zero lateral range and a questionable shape near zero where it is almost flat (i.e., nearly 100%) for a non-trivial distance either side of the search platform's track.

It is impossible to produce a search plan where the tracks are perfectly straight, parallel, and equally spaced relative to all possible search object drift trajectories. Plotting a representative sample of all possible relative search patterns stemming from operationally feasible search plans for a single SAR search situation would likely produce a result that looks very much like "random" search. In all of the U.S. Coast Guard R&DC detection

studies, the search objects were anchored and accurate navigational aids for the participating SRUs were available in the test area, thus removing both systematic and random motion issues.

R&DC studies did find some significant departures from Koopman's model, especially for small search objects combined with poor search conditions. However, the impact of these departures, which were manifested in the form of lateral range curves that were clearly different from that of Koopman's inverse cube model, was not addressed.

Although search craft navigational errors have now been virtually eliminated by GPS, they were a significant additional source of random search pattern error in earlier years.

The "random" search detection function is independent of the exact nature of the instantaneous detection function and shape of the resulting lateral range curve. The "random" search POD depends on only two things: A reasonably uniform distribution of the searching effort over the area searched and the average density (coverage) of that effort. Assuming an approximately uniform distribution of effort can be assured (usually by using parallel track search patterns), then the "random" search POD depends only on the sweep width (W), the effort ($z = vt$) expended in the searched area, and the size (A) of the searched area.

The above considerations seem to indicate that the inverse cube POD vs. Coverage curve is definitely optimistic when applied to SAR. Although the "random" search curve is generally regarded as a lower bound on the POD for an organized uniform coverage search effort, these considerations seem to indicate that the "random" search curve is very possibly a less biased and more realistic estimator of operational SAR POD values than the inverse cube's parallel sweep POD vs. Coverage curve.

One way to help resolve this issue would be to enhance CASP so that it models the simultaneous motions of its simulated search objects ("replications") and search craft while using the lateral range curves derived from R&DC sweep width experiments to estimate the POD for each replication based on its CPA relative to the search craft for each search leg. This enhancement would provide a much more operationally accurate picture of the true effects of searching for drifting objects. It would also provide a valuable tool for determining whether the "random" search curve is the best estimator to use when search planning and evaluation must be done manually. In addition, it would dramatically illustrate the effects of systematic relative search pattern distortion on different search plans and pattern orientations—something that would be an invaluable aid to search planners.

4.6.3 Post-Search Probability Density Distributions

When searching for an object, it is reasonable to assume that searching a particular area without success implies that the chances for the object having been in that area at the time of the search are reduced while the chances for it having been elsewhere at the time are enhanced. This type of reasoning is called *Bayesian inference* and was first developed by the Reverend Thomas Bayes (1702-1761) who first used probability inductively and established a mathematical basis for probability inference (a means of calculating, from the number of times an event has not occurred, the probability that it will occur in future trials). A simplified version of the procedure

for computing a Bayesian update of a CASP probability map is easily illustrated using Figure 4-9 above and Figures 4-14 and 4-15 below.

Suppose the shaded area in Figures 4-14 and 4-15 is searched under ideal search conditions at a coverage of 1.0. Using Koopman's inverse cube model to estimate POD, we get a value of about 78 percent or 0.78 if we accept the traditional interpretation of the POD vs. Coverage graph. (The actual theoretical value computed by Koopman is closer to 79%.) If we multiply the original cell probabilities in the shaded area (from Figure 4-9) by 1 - POD or 0.22, we obtain the values shown in Figure 4-14. This produces an un-normalized probability map. The sum of the cell probabilities in Figure 4-9 was 100 percent, while the sum of the cell probabilities in Figure 4-14 is only 59.74 percent. To restore the sum of the cell probabilities to 100 percent, i.e., to re-normalize the probability map, **all** of the cell probabilities in Figure 4-15 are divided by 0.5974, producing the map shown in Figure 4-15. It is also worth noting that the probability of success (POS) for the search of the shaded area was 1- 0.5974 or about 40.26 percent.

The actual CASP process is more detailed than the above example indicates. Each CASP "replication" is a simulated search object that represents one possible combination of incident location, incident time, search object type, and drift trajectory. CASP can currently have up to 20,000 replications representing a particular situation. In addition to its type, current position and time, and certain other data, each replication has a *P-fail* value. This value is the probability that the object represented by the replication would still be undetected by all searching done to date. The initial *P-fail* value for all replications is 1.0 (100%) since all are undetected prior to the commencement of search operations. For each search where a replication is contained in a search area, its *P-fail* value is updated in an un-normalized fashion just as in the first step of the above example for updating cell probabilities. When CASP is asked to produce a probability map, the *P-fail* values of all the replications contained in each cell are summed and then that sum is divided by the total sum of all *P-fail* values in the simulation to produce the cell probabilities for a normalized probability map.

The allocation of search effort implied by the selection of search area and coverage in this example was not optimal. However, algorithms exist for computing optimal effort allocations for probability maps like those CASP produces. For static distributions, the Charnes-Cooper [1958] algorithm provides a good starting point that can often be transformed into an operationally feasible search plan using appropriate heuristics. Brown's [1980] algorithm and extensions to it may be applied to dynamic distributions.

0.16%	0.33%	3.25%	0.97%	0.42%
0.51%	1.45%	**1.64%**	**1.18%**	**0.73%**
2.43%	3.92%	**1.71%**	**1.65%**	**1.13%**
3.48%	4.43%	**0.72%**	**1.38%**	**1.21%**
2.12%	6.07%	3.33%	2.00%	1.52%
1.40%	3.88%	2.92%	2.05%	1.75%

Figure 4-14. Un-normalized Update of CASP Probability Map.

0.27%	0.55%	5.44%	1.62%	0.70%
0.85%	2.43%	**2.75%**	**1.98%**	**1.22%**
4.07%	6.56%	**2.86%**	**2.76%**	**1.89%**
5.83%	7.42%	**1.21%**	**1.21%**	**2.03%**
3.55%	10.16%	5.57%	3.35%	2.54%
2.34%	6.49%	4.89%	3.43%	2.93%

Figure 4-15. Re-normalized Update of CASP Probability Map.

In contrast, the CSPM is based entirely on a single assumed type of probability density distribution—circular normal. The CSPM does not produce nor does it work with probability maps explicitly, although maps consistent with the CSPM assumptions could be constructed.

The CSPM also does not explicitly produce POS values from which search effectiveness can be judged. This too, could be computed, given the necessary knowledge and statistical tables. However, the CSPM does produce near-optimal search plans within the constraints of limited effort and uniform coverage, but it does so for only one specific set of conditions:

The probability density distribution of search object locations is circular normal and centered on datum,

Either the distribution does not move during the search or the time required to complete the search is short enough to make search object drift during the search negligible,

The effective sweep width is the same everywhere,

The available effort is exactly equal to that required to search the CSPM-recommended square search area (based on the CSPM "safety factors") at a coverage of 1.0.

The CSPM also assumes that as the probable error about the datum position increases with time adrift, the completed square search areas also expand at the same rate with no probability "bleeding" across search area boundaries between searches due to search object motion. In effect, a 10 NM × 10 NM area searched on one day could be treated as a 12 NM × 12 NM searched area when planning the next day's search just as a result of the increase in the total probable error of the datum position after another day of drifting. It was on the basis of this assumption that the CSPM developers justified extending their use of Bayes' theorem when computing the so-called "safety factors" for a static circular normal distribution to distributions that moved and expanded between searches.

Although CASP can generate generalized distributions based on high-resolution environmental data files for computing drift motion, it still suffers from some limitations.

CASP also assumes that the distribution does not move during the search. A "snapshot" of the distribution is taken at the mid-search time and the average POD for the search area is applied simultaneously to all replications contained in the area at that instant.

CASP can generate near-optimal search plans for generalized probability density distributions within certain constraints, but these plans are not always operationally feasible.

Although CASP computes un-normalized Bayesian updates of each replication's *P-fail* value, it uses the same "cookie cutter" approach as used in the above example with cell probabilities. That is, the same POD is applied to every replication in the searched area regardless of how closely or how distantly the search facility may have approached the object represented by the replication at CPA. Since CASP "freezes" the distribution during search updates, the relative motion between search facilities and search objects is not simulated nor are CPA distances computed. Furthermore, no POD is applied to replications immediately outside the designated search area even when they are within the search facility's detection envelope.

CHAPTER 5.

MIN/MAX MODIFICATIONS TO CLASSICAL SEARCH PLANNING

5.1 INTRODUCTION

The entire classical search planning method developed circa 1957 was based on obtaining an estimate of a single mean, or expected, position (***datum***) for the search object together with an estimate of the probable error about that position relative to the search craft. This estimate, called the ***total probable error of position***, was then used to determine the recommended search area. About ten years later, a technique called "minimax" was introduced to handle situations where one of the drift variables, such as leeway rate, could take on one of two significantly different values, and perhaps all intermediate values as well. As we shall see, integrating a method that computed two "extreme" results with one that was originally built to handle only a single "mean" result had a great potential for causing serious problems.

5.2 THE "MINIMAX" TECHNIQUE

Over the years, the Min/Max or "minimax" technique has been applied to several different elements of drift estimation. These have included minimum and maximum leeway rates (e.g., life raft with and without a drogue deployed), leftmost and rightmost leeway divergence off the downwind direction, minimum and maximum distress incident times, minimum and maximum altitude for parachute opening, and even "minimum" and "maximum" distress incident positions (e.g., minimum and maximum distance of a flare from the observer along a line of bearing). Although every drift parameter can have a range of values from some minimum to some maximum, the minimax technique was generally applied to only one variable at a time. The instructions in the most recent edition of the *National SAR Manual* [1991] state, "The SMC [SAR Mission Coordinator] should select the variable with the greatest impact on drift and solve for datum using the possible extremes…" We will examine the instructions for minimax usage in due course. Before doing that, we will look at how the minimax technique evolved.

5.2.1 LEEWAY RATE

In the 1957 *USCG SAR Manual*, a graph was provided for estimating leeway. The graph had two curves, one for a raft with a drogue or sea anchor deployed and one for a raft with no drogue deployed. For leeway with a drogue deployed, the curve was nearly linear at about 2.5 percent of the wind speed. Leeway behavior without a drogue deployed was, according to this graph, decidedly non-linear with respect to wind speed. It ranged from 12.5 percent of the wind speed in very light breezes down to about 3.6 percent of a wind blowing at 30 knots. The maximum difference in the two leeway rates occurred at about 14 knots when it was just under 0.6 knots (0.958 – 0.375 = 0.583). The leeway direction was assumed to be directly downwind in all cases. Min/Max solutions were not mentioned.

In 1967, Amendment 4 to what in 1959 had become the *National SAR Manual*, contained the first Min/Max procedure. This procedure was applied only to life raft drogue/no drogue leeway rates. The method instructed the search planner to compute two "datum" positions, $datum_{min}$ and $datum_{max}$, based on the "minimum" leeway rate (drogue deployed) and the maximum leeway rate (drogue not deployed). A circle was then drawn around each position with a radius equal to one-eighth of the respective distances drifted. Finally a single large circle was drawn about a point between the two "datums" such that the large circle contained and was tangent to the two smaller circles. The center of this large circle was called "$\text{datum}_{\text{minimax}}$" and its radius was taken to be the probable drift error and called "$d_{e\ \text{minimax}}$." Figure 5-1 illustrates the results of the Min/Max procedure applied to the life raft problem when it is not known whether a drogue has been deployed. The leeway rates used were from the original leeway graph for a mean wind of 240°T/14 knots. The total water current was assumed to be 000°T/1.0 knot and the time adrift was taken to be 24 hours. The drift error in this example was estimated as one-eighth of the distance from the starting position (0,0) to be consistent with the doctrine that was in effect at the time this method was introduced.

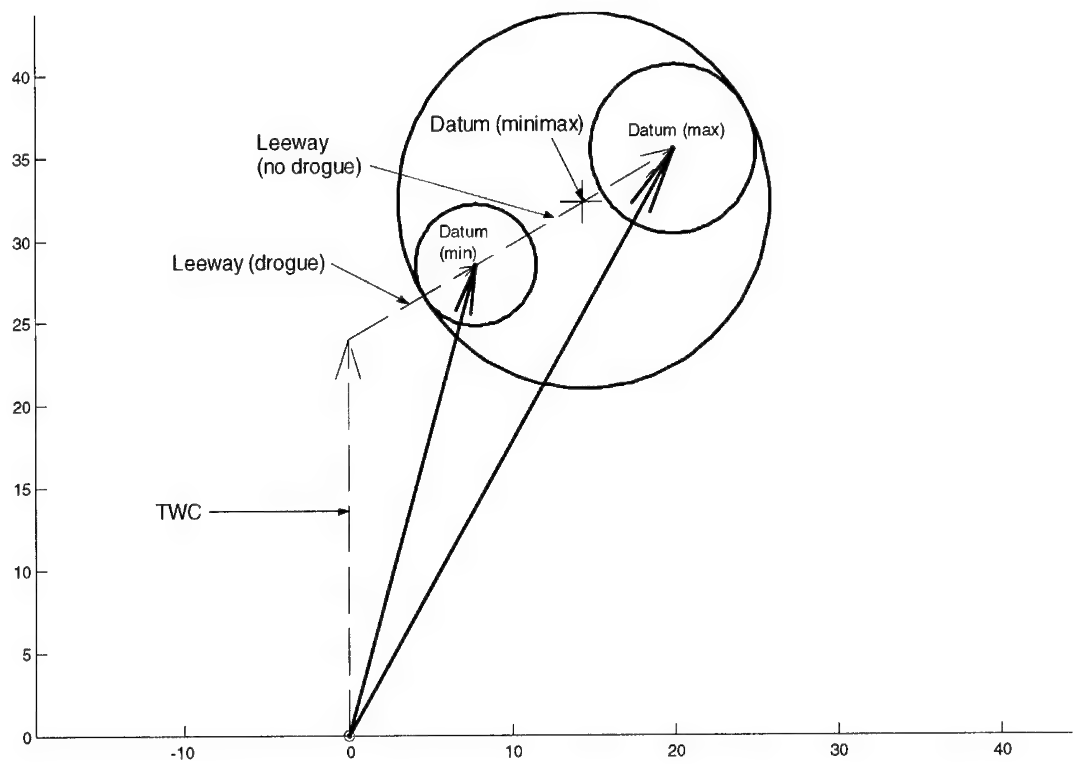

Figure 5-1. Min/Max Applied to Leeway Rate.
(D_e = 0.125 × distance drifted)

It is instructive to study Figure 5-1 carefully. Prior to the introduction of Min/Max solutions, each of the smaller circles would have been taken to be the probable error in the search object's estimated position (ignoring incident and search craft position errors for the moment) for the drogue and no-drogue scenarios, respectively. It would have been assumed that the distributions of possible locations around each "datum" were circular normal. Under these assumptions, the union of the two circles would contain about 50 percent of the search object's possible locations.

However, this means the large circle containing both of the smaller circles cannot possibly be the 50 percent containment contour on a circular normal distribution nor even a close approximation. Nevertheless, its radius, $d_{e\ minimax}$, is used as the probable drift error when it is combined with the incident and search craft position errors to compute the total probable error of position, E. Furthermore, the same "safety factors" are used with this greatly increased estimate of E to compute the recommended search area.

Clearly the large circle in Figure 5-1 must contain much more than 50 percent of the search object's possible locations under the assumptions given. Furthermore, the distribution is not circular normal and the large circle must contain significant amounts of area where the search object has little or no chance of being located. This will become increasingly true as the "safety factors" increase with continued searching. In short, there is a great potential for wasting scarce resources searching in very unlikely places if Figure 5-1 is used to plan a search.

In 1986, the "confidence factor" for estimating drift error was increased from one-eighth to three-tenths of the distance drifted. Figure 5-2 shows the impact of this change.

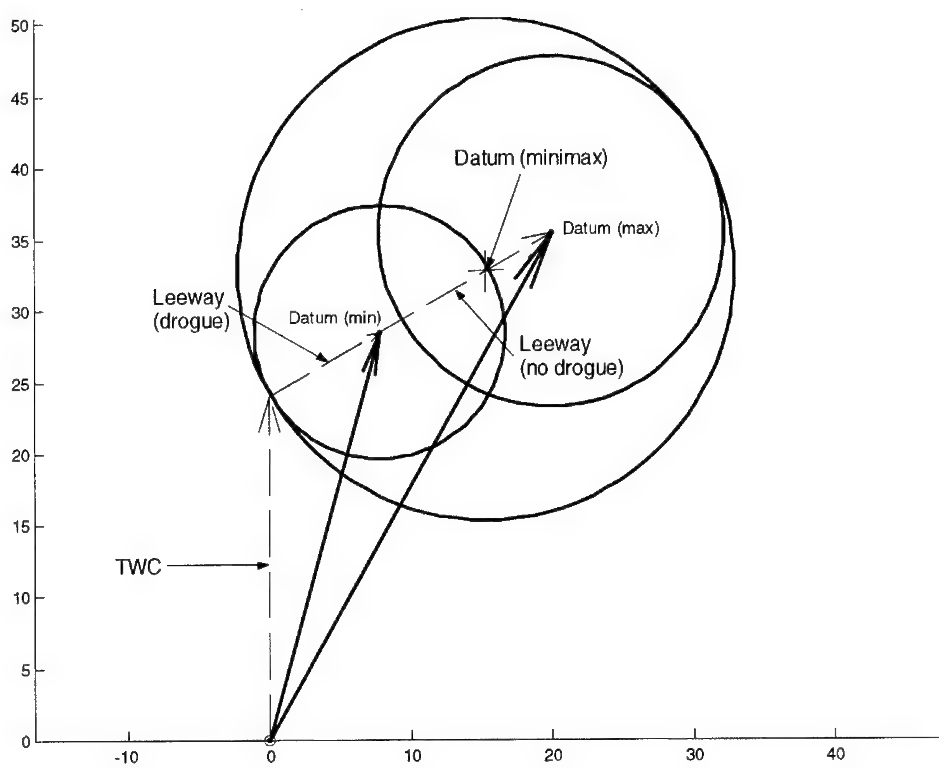

Figure 5-2. Min/Max Applied to Leeway Rate.
($D_e = 0.3 \times$ distance drifted)

This increase in the drift error estimate was based on some comparisons between actual trajectories of satellite-tracked drifting objects and predictions made by using the methods prescribed for search planning. It was concluded that the drift error factor of one-eighth significantly underestimated the probable error in SAR drift predictions and that three-tenths was a more realistic figure.

It is instructive to analyze this solution to the same problem as that of Figure 5-1, the only difference being the different estimates of the drift distance error. Note that although the large circle in this case is larger than in the previous example, the difference between its area and the combined areas of the two smaller (but larger than before and now overlapping) circles is not nearly as pronounced. In other words, if the larger estimate of D_e is more appropriate, then approximating the two probable error circles with a single circle drawn in the prescribed fashion may not be as unreasonable as it previously appeared.

Finally, suppose we estimate the total probable error in the drift velocity estimate, DV_e, not as a percentage of the mean drift velocity, DV, but as a function of the quality of the inputs used to estimate DV. If there were four inputs to the computation of DV and each had a probable error of 0.3 knots, then by Theorem 3 of section 4.2.3, DV_e would have been 0.6 knots. If we multiply this value by the 24 hours the object was adrift, we get the results shown in Figure 5-3 below.

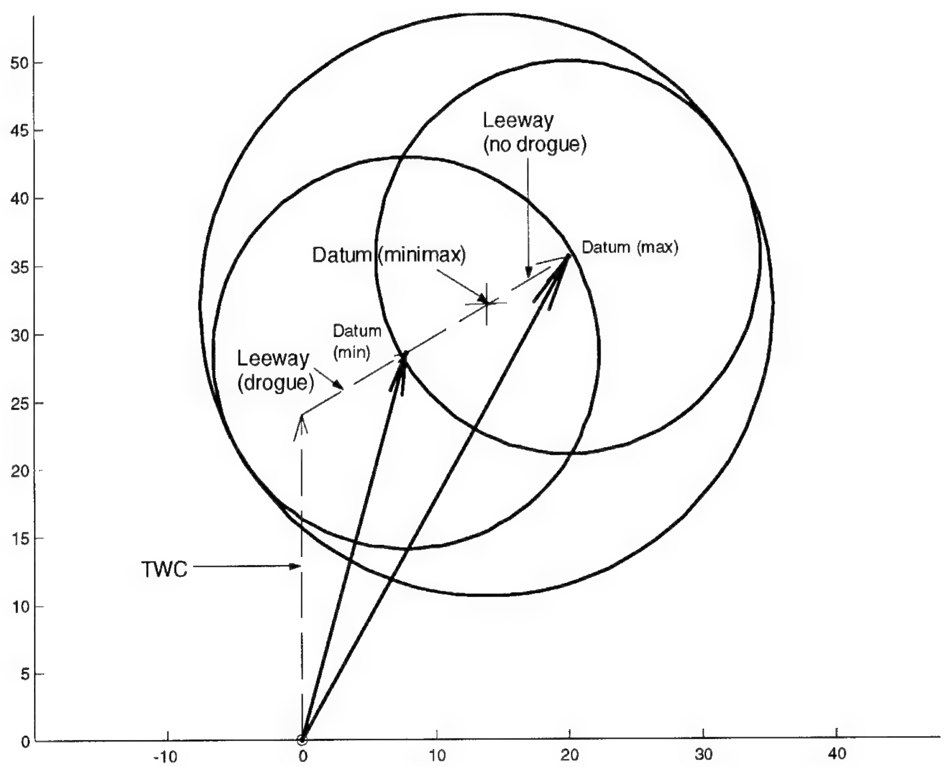

Figure 5-3. Min/Max Applied to Leeway Rate.
(DV_e = 0.6 knots)

Note that both of the smaller circles now have the same radius as a consequence of computing drift distance error on the basis of the probable error in the drift velocity estimate and the amount of time the object has been adrift rather than as the same percentage of two different drift distances. The increase this change caused in the estimate of the "maximum" probable error is relatively small while that of the "minimum" probable error increased much more substantially. If this estimate of the "minimum" and "maximum" probable position errors is valid, then the single large circle approximation continues to improve.

5.2.2 Integrating Min/Max with Classical Search Planning

The most obvious explanation for why the originators of the Min/Max technique chose to establish a single "$datum_{minimax}$" as the center of a single large circle is that they wanted to integrate the Min/Max results into the basic classical methodology. Recall that the classical search planning method is based on establishing a single datum and a probable error about that datum. However the method is also heavily dependent on the assumption of a circular normal distribution—an assumption that the Min/Max results just described clearly violated for the small "minimum" and "maximum" drift error estimates being made at the time, as Figure 5-1 shows. It seems likely that the originators of the Min/Max technique acting some ten years after the classical method was developed were unaware of the underlying statistical assumptions on which that method was based.

It seems obvious that in all of the situations represented by the examples given in paragraph 5.2.1, a more nearly optimal allocation of the available search effort could be obtained by covering a rectangular area centered on the line connecting the "minimum" and "maximum" datums and including both of them. This is especially true if intermediate leeway rates are also possible. For the life raft problem we have been dealing with, intermediate leeway rates were clearly possible. First, there was no guarantee that either a drogue was always deployed or it was never deployed. Second, there was also no guarantee that even if it was always deployed that it was deployed correctly so as to achieve the maximum effect. Either situation would produce, in effect, an intermediate leeway rate. In allowing for such possibilities, it would be reasonable to assume that a ridge of high probability density exists between the "minimum" and "maximum" datums as depicted in Figure 5-4 below.

However, there are at least two things to consider. First, it would be difficult to establish an easy manual method for determining the size and placement of the optimal rectangle, although establishing the optimal search radius for one datum, drawing circles of that radius around both datums and circumscribing a rectangle around the two circles would produce a near-optimal result. Second, the situation is not static. Life rafts continue to drift during the search. One of the most effective ways to deal with this issue is to orient the search legs parallel to the mean direction of motion (which is usually different from the leeway direction due to the presence of currents). A single square circumscribed around a circle has the advantage that it can be oriented in any direction that is desired. Finally, if the drift errors are as large as now believed, then perhaps at the end of the day the Min/Max procedure did approximately the right thing, even if it was for the wrong reasons, at least when dealing with life rafts where the drogue's status (deployed or not deployed) was unknown.

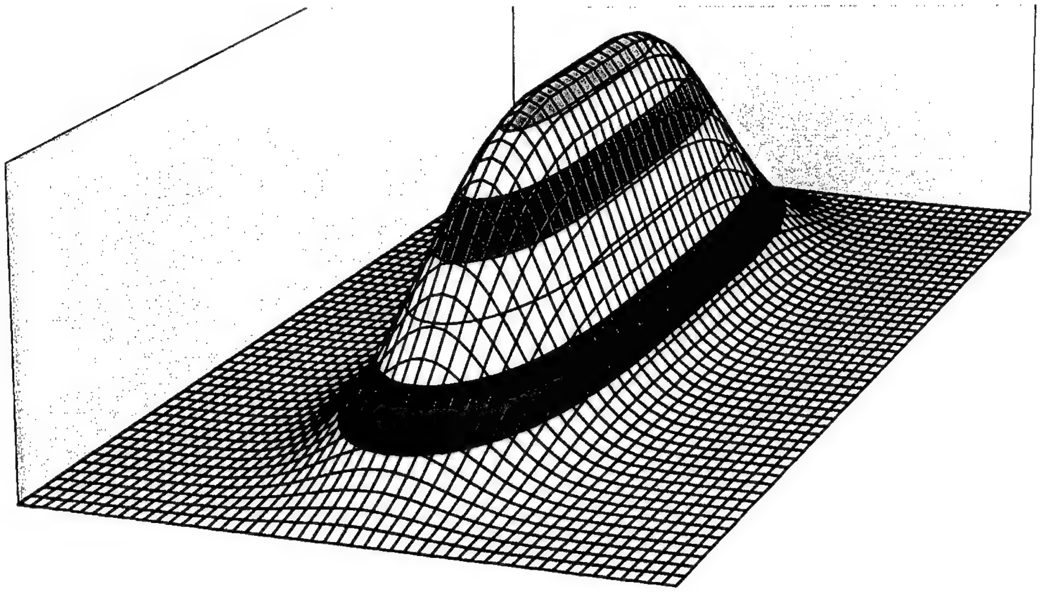

Figure 5-4. Ridge of High Probability Density Between $Datum_{min}$ and $Datum_{max}$.

5.2.3 Leeway Divergence

Leeway divergence was first mentioned in a SAR context in Amendment 3 (1963) to the *National SAR Manual*. A statement was made that a search object's leeway could be "up to 40 degrees" off the downwind direction. The corresponding instructions directed search planners to continue computing leeway in the downwind direction, but to consider expanding the search area to the left and right of the downwind direction to account for divergence. No explicit instructions for determining the dimensions of such an expanded search area were provided, however. The only source of leeway rates remained the original life raft leeway graph.

Amendment 8 in 1972 provided leeway rates as a percentage of the wind for objects other than life rafts. The source of these values is not clear, but they were probably preliminary results from R&DC experiments reported by Hufford and Broida [1974]. No additional divergence information was provided in Amendment 8.

The *National SAR Manual* was completely rewritten in 1973. Amendment 2 in 1976 introduced a new leeway graph derived from Hufford and Broida [1974] that covered a wide range of objects, including life rafts. The previous life raft leeway graph was removed. Three maximum leeway divergence values were also provided: ±35 degrees for rubber rafts, ±45 degrees for craft with moderate to deep keels, and ±60 degrees for craft with relatively shallow draft. No information was provided about the distribution of divergence angles except to say that the chances for an object's leeway diverging to the right of the downwind direction was equal to the chances it would diverge to the left of the downwind direction. Generally this was interpreted to mean that any divergence angle between the two extremes was equally likely. Despite this added information, the instructions in the *National SAR Manual* still called for computing leeway in the downwind direction only. The manual was again rewritten in 1991. The drift error factor had been increased from 0.125 (1/8) to 0.3 in 1986, but no changes related to leeway information or

instructions were made. However, this was not true of the *USCG Addendum to the National SAR Manual.*

Change 1 to the *USCG Addendum to the National SAR Manual* in 1991 provided job aids and instructions for computing datums. Three types of "Min/Max" solutions were included on the Leeway Worksheet, and one of these was "directional uncertainty" for dealing with leeway divergence. For the first drift interval, two leeway vectors were computed, one for the maximum leeway divergence to the left of the downwind direction and one for the maximum leeway divergence to the right of the downwind direction. These leeway vectors were then used to compute two datums. Probable error circles were drawn around each datum (using 0.3 times the distance drifted for the probable error radius) and a single large circle centered between the datums was drawn around them in the same fashion as for the earlier Min/Max leeway rate solution. Figure 5-5 illustrates a solution for an object with a maximum leeway divergence of ±60 degrees and a leeway rate of 5 percent of the wind speed when the wind is 240 degrees T/20 knots and the TWC is 045 degrees T/1.0 knot.

Figure 5-5 has some interesting characteristics when compared with the earlier scenarios involving the unknown status of a raft's drogue. Perhaps the most striking features are the distance between the two datums and the consequent very large size of the "minimax" drift error circle. Recall that the maximum separation rate between rafts with and without drogues deployed was about 0.6 knots. In the example just given, the "minimum" and "maximum" datums are separating at more than 1.73 knots. This is why the error circles around the "minimum" and "maximum" datums do not overlap despite the use of a larger drift error factor (0.3). Also recall that in the case of rafts with or without drogues, the "minimum" leeway rate was not really a minimum but the mean rate for a raft with a drogue deployed. Similarly, the "maximum" leeway rate was not really a maximum but the mean rate for a raft without a drogue deployed. In contrast, the values given for leeway divergence angles were represented as true extreme values relative to the downwind direction. Even a cursory inspection of Figure 5-5 reveals that use of this method implied that even larger divergences off the downwind direction were not only possible but would be present in a significant number of cases. Finally, the size of the circle centered on $datum_{minimax}$ and its use in determining the recommended search area produces several additional illogical implications.

Recall that, for search area determination purposes, the Min/Max technique treats the radius of the large circle, $D_{e\ minimax}$, as if it is the probable error of a circular normal distribution of possible search object locations. (We will ignore incident and search craft position errors for the moment for simplicity. Their inclusion would only exacerbate the problems about to be raised.) If the $datum_{minimax}$ position is the center of the distribution, then it represents the mean drift. In this example, this implies that the mean drift rate was 060 degrees T/1.5 knots and had a probable error of about 1.39 knots, which is about 92 percent of the mean value. The "minimax" circle is so large it nearly includes the incident position. If it does represent the 50 percent containment contour on a circular normal distribution, then there is a small, but still significant, probability that the search object drifted upwind and up current. Since the wind and current are only mean values, this can be a valid implication if the uncertainties about those mean values are large enough. However, the solutions for both drift trajectories assumed the effects of drift error amounted to only about 30 percent of the mean drift speeds, not 92 percent.

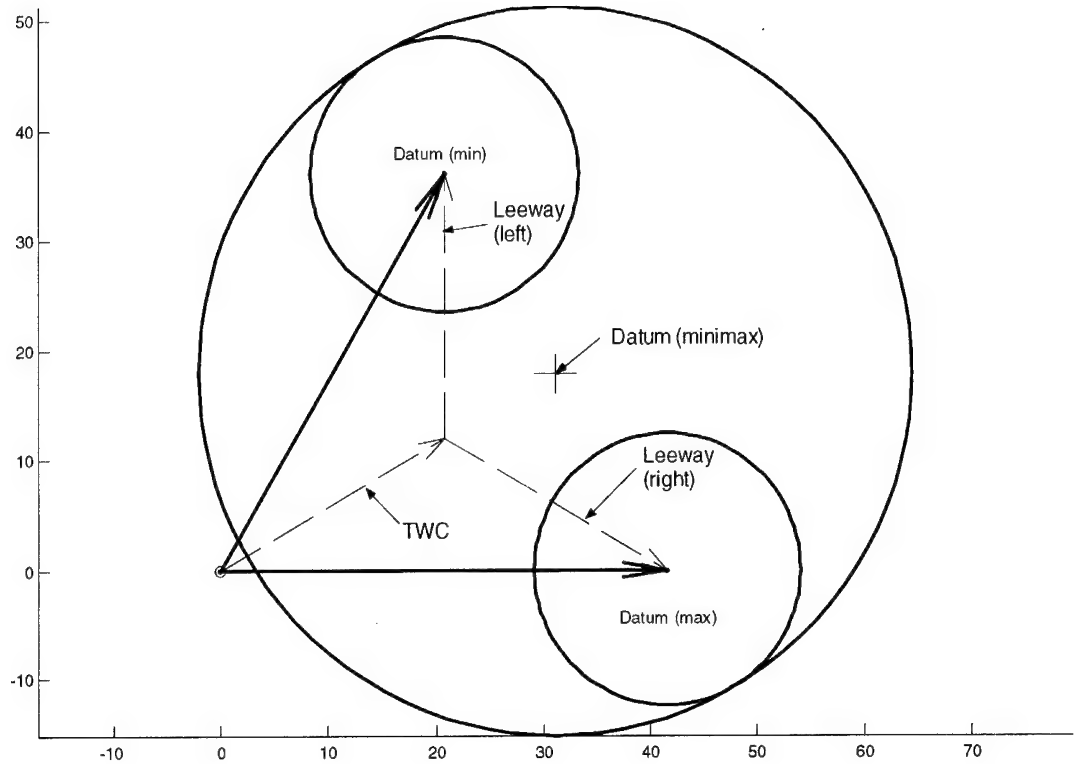

Figure 5-5. Min/Max Applied to Leeway Divergence.
(D_e = 0.3 × distance drifted)

The large circle was merely trying to address the uncertainty about which side of the downwind direction the leeway was on, i.e., the uncertainty about which "tack" the search object took. Clearly, it took in much more area than necessary to meet this need. The Canadian computer program CANSARP addresses this issue by computing datums for the maximum left and right divergence angles just as in Figure 5-5 but it also computes datums (with error circles for each) for nine intermediate divergence angle values, producing an "arc of probability" that covers a much smaller area. Although CANSARP continues to compute and display a single large circle, in practice, search areas are often formed by circumscribing rectangles around the "arcs" formed by the eleven datums and their respective error circles. This may be a more reasonable approach, but it is still inconsistent with the latest and best available information on leeway behavior.

Before we leave our examination of Figure 5-5, two other items are worth considering. The first is the ambiguity of trying to apply Min/Max to vector quantities. Note that the datum associated with the "minimum" leeway direction (000°T) is farther from the incident position than the datum associated with the "maximum" leeway direction (120°T). In reality, neither of these datums represents a "minimum" or "maximum" possible distance from the incident position. The second issue to consider is how to proceed with a second drift interval. Should each of the datums shown in Figure 5-5 be treated as the starting point for another pair of "minimum" and "maximum" datums for the second drift interval? Should the object with leeway to the left of the downwind direction continue to the left of the new downwind direction and similarly for the object with leeway to the right of the downwind direction, producing only two new datums? Or

should all Min/Max computations always start from the incident position, thus producing only two datums? Neither the *National SAR Manual* nor the *USCG Addendum* offered any guidance on how to proceed.

5.3 QUESTIONS ABOUT THE LOGIC OF MIN/MAX

Use of the Min/Max technique has generated many questions and much confusion over the years, almost all of which have remained unresolved. Clearly the intent of the Min/Max technique was to provide search planners with guidance for situations that did not seem to be addressable by the classical search planning method. Whether the guidance provided was good or poor is an issue worth investigating. We have seen two applications of Min/Max, one of which may have produced a somewhat reasonable search area while the other clearly produced an excessively large search area.

Although the general intent of Min/Max seems clear enough, it is difficult to state its intended results any quantifiable terms. Even the *National SAR Manual* is inconsistent in its summary of Min/Max results. Paragraph 516 C of the *National SAR Manual* states,

> "The SMC should select the variable with the greatest impact on drift and solve for datum using the possible extremes, such as the faster speed of an unballasted raft and the slower speed of a half-swamped boat. This establishes the maximum and minimum drifts. Datum minimax is half way between these points, ensuring that the most probable position is closest to the center of the search area."

We have just seen that these instructions and assertions are inconsistent with the actual implementation of the Min/Max method.

The term "greatest impact" is ambiguous. The variable with the "greatest impact" could be the one that makes the largest mean contribution to drift speed. Often the wind, which causes both leeway and wind current, would have the "greatest impact" under this interpretation. Another interpretation would be the variable that had the greatest difference between its maximum and minimum magnitudes. That variable would also be a candidate. The variable whose extreme values, when combined with the mean values of the other variables, produced the greatest difference in drift speed (or direction) could also be considered a candidate.

It is impossible for a search planner to use "the possible extremes" since that information is not provided. Leeway graphs, for example, give the search planner an estimate of the mean leeway for objects such as unballasted rafts and provide no information on half-swamped boats.

Min/Max does not establish the minimum and maximum drifts. First of all, data on "extreme" values is unavailable. Second, even if they were available, combining the extremes of one variable with the means of all the others does not guarantee extreme results.

There is also no guarantee that the point midway between the two computed datums will be the place where the probability density peaks. In the example used in the first sentence of paragraph 516 C quoted above, two such dissimilar objects would quickly become widely separated with no possibility of intermediate leeway rates. This means that intermediate drift rates are unlikely

and that the survivors are much more likely to be found near one datum or the other than at a point half way between them.

Finally, effort allocation depends on more than just probability density. In the case of two dissimilar objects, the sweep widths are also likely to be different, making an optimal search plan for one sub-optimal for the other. In the example cited, it would be better to plan separate search efforts around each datum since it is likely that two mutually exclusive scenarios are involved.

The following list of issues demonstrate some of the serious problems with Min/Max "logic" as it has been applied to search planning.

1. Min/Max does not live up to its name. For example, the leeway rate provided for a life raft with drogue deployed is a mean value, not an extreme value. That is one reason why a probable error circle is drawn around the "minimum" datum rather than using the "minimum" datum itself as a limiting value. Similar logic applies to the "maximum" datum.

2. Application of Min/Max does not produce the minimum and maximum distances the search object could have drifted.

3. It is not always clear how Min/Max should be applied to several drift intervals in succession. Given a $datum_{min}$, a $datum_{max}$, and a $datum_{minimax}$ from the first drift interval, the search planner is faced with several choices for computing datums for the second interval.

   a. The search planner could compute a new "minimum" and "maximum" datum pair using each of the previous datums (min, max and minimax) as a starting point. This would result in six datums and raise questions about where a new $datum_{minimax}$ should be placed and how the radius of a single all-encompassing circle should be computed.

   b. The search planner could also compute only two new datums—a "minimum" datum using the previous "minimum" datum as a starting point and a "maximum" datum using the previous "maximum" datum as a starting point. Depending on how the mean wind and current over the second interval relate to those of the first interval, especially in terms of direction, the distance between the "minimum" and "maximum" datums could either increase or decrease. A decrease could cause the radius of the large circle, $D_{e\ minimax}$, to decrease as well, the implication being that the uncertainty about the object's location at the end of the second interval is less than it was at the end of the first interval. This is generally an illogical result for objects adrift in the marine environment.

   c. Finally, the search planner could compute a new Min/Max solution by using the incident position for the starting point and the mean wind and current over both intervals for computing the "minimum" and "maximum" drift trajectories.

4. All of the variables used in drift computations have some uncertainty associated with them. The amount of uncertainty can be expressed in several ways, including standard deviation (standard error), probable error, or as minimum and maximum values. From time to time it has been suggested that datums be computed based on all possible combinations of minimum and maximum values for the drift variables and that a complete set of new datums be computed for every previously computed datum as time progresses from one drift interval to the next. The apparent goal of such a scheme is to "bound the problem," i.e., determine the

smallest geographic area that is guaranteed to contain the object. This has several drawbacks. First, it does not guarantee search object containment. There may well be combinations of values between the "minimum" and "maximum" values of the drift variables that produce datums outside the area indicated. Second, it leads to a "combinatorial explosion" of datum computations that would quickly overwhelm even a quite large computer. Third, it would provide no information on the distribution of search object location probability density and would therefore provide no information on how to optimally allocate the available effort within the bounded area.

5. Most of the variables used in drift computations are two-dimensional vectors with uncertainties (probable errors). Vectors may be expressed as an ordered pair of orthogonal components or as a direction and magnitude. This means vectors are "bivariate" and so are their uncertainties. There is no logical way to define a "minimum" or "maximum" value for a vector like there is for a scalar quantity. In the case of the life raft with unknown drogue status, this problem is avoided by holding leeway direction constant (directly downwind) and examining the mean minimum and mean maximum leeway rates which are scalars. In the case of leeway divergence the leeway rate relative to the mean wind is held constant and the mean "minimum" (leftmost) and "maximum" (rightmost) leeway directions relative to the mean downwind direction are examined.

In summary, the Min/Max technique, as implemented, seems to have been the ill-fated result of applying one-dimensional scalar reasoning to two-dimensional vector quantities, exacerbated by a lack of knowledge about the statistical and search theory basis of the classical search planning method.

CHAPTER 6.

COMPARISON OF SEARCH PLANNING AND CASE MANAGEMENT TOOLS

6.1 INTRODUCTION

Search planning tools may be grouped into three broad categories. These are:

Manual methods

Automated versions of manual methods

Sophisticated stochastic computer models or simulations

Neither automated search planning nor case management tools were provided for evaluation with the exception of the Coast Guard's C2PC/AMS search planning tool. However, another search planning tool was obtained for evaluation and brief demonstrations of two others were obtained. Brief demonstrations of two case management tools were also obtained. One of these had already been evaluated by the Coast Guard's Atlantic Area Command Center.

6.2 MANUAL SEARCH PLANNING METHODS

All manual methods may be traced back to the Classical Search Planning Method (CSPM) discussed earlier. We speak of "methods" in plural because a number of significant modifications to the CSPM have been made over the years. These modifications appear to have been "field changes" initiated, developed and implemented without the careful research that went into the CSPM's initial development. The most important of these modifications was the "minimax" technique discussed above.

In 1992, the International Maritime Organization (IMO) and the International Civil Aviation Organization (ICAO) established a joint working group for the harmonization of aeronautical and maritime SAR. Many of the issues were organizational. For example, conventions of each organization established an international requirement for Rescue Coordination Centers (RCCs) and delineated areas of responsibility to cover all international waters and the airspace over them. However, the areas of responsibility for maritime SAR and aeronautical SAR did not always coincide. It was possible for two RCCs in two different countries to be responsible for SAR in the same region—one for aeronautical cases and one for maritime cases. There was also a need for specific protocols and standard communications techniques to be established among RCCs to coordinate responses to cases near boundaries and responses to large-scale incidents where assets from a number of countries might be required. Finally, there was a need to standardize search planning and coordination methods.

After several years of deliberations, review of many SAR-related documents, and some additional research, the *International Aeronautical and Maritime Search and Rescue (IAMSAR)*

Manual was published in 1999. This manual contains an updated version of the CSPM that explains and uses the concepts of probability of containment (POC) and probability of success (POS) that were hidden from the search planner's view in the CSPM. Instead of fixed "safety factors" and "mid-point compromise" techniques for allocating effort to a "point datum" (circular normal distribution), the *IAMSAR* method replaces these with continuous optimal search factor curves that allow the search planner to correctly match the search area size and coverage to the level of searching effort available. This capability was extended to line datums and a "trial-and-error" approach to a generalized probability density distribution was also presented. Although they are too cumbersome for manual use in maritime SAR, probability maps are presented as an aid to understanding the optimal effort allocation process and the proper use of previous negative search results in allocating effort for the next search. For problems over land, such as searching for a downed aircraft, probability maps can be used and kept properly updated by hand since search object drift and its associated uncertainties are not an issue.

The *IAMSAR* method initially abandoned the minimax technique to avoid all the ambiguities and confusion of that method. There were also preliminary indications from recent leeway experiment results that divergence angles were not nearly as large, on average, as the maximum values previously provided to search planners, making leeway divergence appear to be less of an issue. As the leeway experiments progressed, however, it became apparent that although leeway divergence angles were smaller on average than previously thought, they were still significant. As a result, a method for dealing with leeway divergence based on the research of Allen and Plourde [1999] was developed for the *IAMSAR Manual*. Unlike its minimax predecessor, this method uses mean, not extreme, divergence angles and is properly integrated with the remainder of the manual search planning process. This includes near-optimal search effort allocation that is a substantial improvement over the minimax technique, leading to more efficient use of assets, shortened searches on average, and more lives saved. However, it is far from being a panacea. The limitations imposed by keeping the method within the paper-and-pencil realm force it to maintain the grossly oversimplified approach of the CSPM, especially in the maritime arena.

The original CSPM, of course, was developed long before significant computing power and automation were readily available. It was a truly manual, pencil-and-paper, technique. Over the years, some modifications were made to the CSPM, such as the minimax technique discussed in the last chapter and its extension to quantities such as leeway divergence. Other modifications included more complex ways to compute wind current (which have since been discarded), improved leeway graphs that are based on significant amounts of field research and cover a wider variety of objects, improved sweep width tables for unaided visual search and additional tables for visual aids (e.g., night vision goggles (NVG)) and other sensors (e.g., forward-looking infrared radar (FLIR)), also based on significant amounts of field research by the USCG Research and Development Center. Methods for dealing with additional environmental variables, such as tidal and rotary currents, were also added. The question of how to allocate search effort when the available effort was significantly different from that required to complete a coverage 1.0 search of the recommended search area was addressed by a technique called "mid-point compromise." Finally, the amount and resolution of the available environmental data increased dramatically, making it necessary for search planners to compute more datums at more frequent intervals if they wanted to use this data to best advantage. Each of these changes typically added some computational or other complexity so that by the time computer support

did become readily available, the computational burden placed on the search planners had become substantial. This led to the next category of search planning tools.

6.3 AUTOMATED MANUAL METHODS

The primary purpose of computerized or automated versions of manual methods is to provide the search planner relief from much of the computational burden that had been imposed on the manual methods over the years. The more modern automated tools also relieve the search planner of much or all of the plotting burden associated with search planning on paper charts. Modern Geographic Information Systems (GIS) provide a great many conveniences for the search planner. For example, a search planner can accurately display a chart on the computer screen and simply point and click on a chart location to have the latitude and longitude of that location automatically entered into the appropriate form on the screen and the appropriate program variables in the search planning software. Search areas are easy to specify and manipulate on the screen, as are the search patterns they contain. Directions and distances between points may be easily found by pointing first to one point and then to a second point with a mouse or other pointing device. As the cursor is moved across the screen, its position on the displayed chart in latitude and longitude is automatically displayed. There are many other useful features as well.

Automated versions of manual methods are generally easy to use, produce results quickly, and provide many conveniences over pencil-and-paper methods. Needless to say, modern versions of these tools are very popular with search planners. However, they have one important drawback. Aside from avoiding some common sources of human error (e.g., computational mistakes, plotting errors, etc.), and, in some cases, making slightly better use of the environmental data now available, these tools do not substantially improve the quality of the resulting search plan over that of truly manual methods. For example, the problem of situations where the initial distribution is not circular normal, if addressed at all, is handled by choosing a small number of incident positions and times, then solving each as an individual search problem. The search planner is either left to try and make sense of the several recommended search areas that result or the software simply circumscribes a rectangle around all of them. Increasingly available high-quality, high-resolution environmental data are not used effectively. None of the automated manual solution tools surveyed improved search effort allocation over the manual methods on which they are based. These tools do not even attempt optimal effort allocation over a single-point datum, something that could be done with relative ease, and in fact has been done for the manual method given in the *International Aeronautical and Maritime Search and Rescue (IAMSAR) Manual* [1999].

Automating the manual methods has also provided additional opportunities to modify them. In fact, automation ***required*** making some modifications. Prior to automation, the manual method of choice the world over was generally a clone of the method given in the U.S. *National SAR Manual* edition current at the time. However, the method as described there has always contained a number of omissions and ambiguities. Search planners were often encouraged to consider certain factors when planning a search, but specific quantitative guidance (e.g., how much to change the location, size, shape, orientation or coverage of search areas) was not provided. The question of how to compute minimax solutions subsequent to the first search was not addressed. Since these and other questions had to be addressed in order to develop useful software tools and aids, software developers and their clients in the SAR community were left to

their own devices to resolve them. Not surprisingly, several different approaches to these and other issues were implemented. The sections that follow contain highlights of the different approaches that have been used by software developers to address certain issues.

6.3.1 USCG Search and Rescue Planning (SARP)

The first automated version of a manual search planning method was the USCG's SARP program developed circa 1970. It was part of a system of programs and data files developed to provide RCC Controllers with operational information for use in the prosecution of SAR cases. Improved Coast Guard SAR performance was the main objective and it was to be accomplished through:

Eliminating the potential for computational errors.

Increasing the time available for gathering and appraising specific case information, due to the time saved by the rapid computational ability of the computer.

The function of the SARP program was to develop a solution to the search planning problem based on the doctrine specified in the National Search and Rescue Manual and the search planning methodology taught at the National Search and Rescue School.

SARP was designed to require only four inputs:

Incident date and time.

Last known position of the distressed craft.

Probable position error of the distressed craft.

Probable position error of the search craft.

SARP was also designed to accept a number of optional inputs to improve its flexibility and the accuracy of the computed solution. SARP's features included:

Sweep width computation based on search object length, percent cloud cover, meteorological visibility, and search craft altitude input by the user.

Local wind current calculation, based on the graph given in James [1966]. The wind current could be fetch-limited, duration-limited, or fully developed based on the user's inputs.

Leeway calculation based on formulas and parameters developed from U.S. Coast Guard Research and Development Center leeway experiments.

Use of surface wind data that could be input by the user or taken from the on-line data files that were updated twice daily with analysis and forecast data received from Fleet Numerical Weather Center (now Fleet Numerical Meteorology and Oceanography Center), Monterey, CA. Wind data input by the user included direction, speed, and time of the "observation."

Use of average sea currents that could be input by the user or obtained from on-line data files based on climatology. The on-line files included most of the northern hemisphere on a one-degree grid and several regional higher-resolution files near the U. S. Atlantic coast. One of these, the Gulf Stream file, was updated on a weekly basis by oceanographers at the International Ice Patrol interpreting a NOAA thermal imagery product. Sea current data input by the user included the position of the "observation," direction, speed, and whether the data was based on DMB observations (in which case it was taken to be total water current data and other sources of current were not used).

Use of winds aloft data input by the user to compute parachute drift for cases involving bailouts.

Computing a DR position for the distress incident time from the distress craft's last known position, the time of that position, and the distressed craft's estimated course and speed.

SARP was accessed from USCG RCCs via the SARLANT and SARPAC polled teletype networks. RCC personnel would complete paper forms by hand and give them to their servicing communications center. The radioman on watch would then produce a teletype "service" message containing the data on the forms in the prescribed format and transmit it to the central computer at the Transportation Computer Center in Washington, DC. The computer output would then be transmitted back to the requesting RCC via its servicing communications center. By the standards of the day, SARP was a very fast program and within a very few minutes the RCC would have the computer's response in hand.

SARP computed drift trajectories on a one-hour time step. At the beginning of each hour of "simulated" drift, both wind data and sea current data were accessed using "nearest data point" logic with respect to the last computed search object position. For user-supplied wind data, the "nearest data point" was computed with respect to the simulated time. For user-supplied sea current, the "nearest data point" was computed with respect to the search object's computed position. For gridded data from the on-line files, the nearest data point in both space and time was used. User-supplied data did not have to be provided at either fixed or uniform intervals of time for wind or at fixed or uniform spatial intervals for current.

A SARP-generated solution contained some intermediate drift positions at regular intervals. The exact frequency of intermediate positions shown depended on the length of the total drift interval. The datum position(s) were also displayed (in two columns for minimax solutions) for either the time at which the computer received the request or as of the time requested by the user. The computed total probable error of position was provided. The standard search radius based on the search number (first, second, third, etc.) and corresponding "safety factor" was also computed and provided. The size of the corresponding square search area in square nautical miles was displayed along with the latitudes and longitudes of its boundaries, based on assuming it was oriented so parallels and meridians formed the sides. If the necessary data had been provided, the computed sweep width was also displayed, along with the times of sunrise and sunset at the datum position. Finally, the average drift vectors over the period and over the last hour were displayed. Figure 6-1 shows a SARP solution as received via teletype.

```
CF CA DE CD
TEST CASE PACAREA
BT
UNCLAS
USCG COMPUTERIZED SAR PLANNING SYSTEM
PROGRAM SARP EXECUTED AT 211500Z JUL 75
        INCIDENT TIME AND LAST KNOWN POSITION DATA
INCIDENT DTG     201300Z JUL 75
LADT KNOW POS    37-30.0N   123-00.0W
SEARCH NUMBER   3   DATUM DTG    221440Z JUL 75 FIRST LIGHT
            DISTRESS CRAFT AND LEEWAY DATA
INITIAL POSITION ERROR (X) 15   SEARCH CRAFT ERROR (Y)  5
TARGET LENGTH, CLOUD COVER, VISIBILITY, SEARCH ALT   2  40  15  1000
LEEWAY CONFIGURATION IS F 60 0.07 1.0
              SURFACE WINDS - MONTEREY
NUMBER OF WINDS IS  10
NO DIR/SPD         DTG      FETCH      NO DIR/SPD         DTG      FETCH
 1 180/10  180600Z JUL 75  999         6 150/17  201800Z JUL 75  999
 2 190/15  181800Z JUL 75  999         7 150/13  210600Z JUL 75  999
 3 220/20  190600Z JUL 75  999         8 150/13  211800Z JUL 75  999
 4 210/20  191800Z JUL 75  999         9 150/13  220600Z JUL 75  999
 5 160/15  200600Z JUL 75  999        10 150/13  221800Z JUL 75  999
 ......................................................................
INTERMEDIATE DRIFT POSITIONS
HOUR   POSITIONS            ASC     POSITIONS            ASC
  0    37-30.0N 123-00.0W
 12    37-37.1N 123-11.3W  NAVO    37-43.3N 122-59.8W  NAVO
 24    37-44.2N 123-22.6W  NAVO    37-56.6N 122-59.6W  NAVO
 36    37-51.2N 123-33.9W  NAVO    38-09.9N 122-59.4W  NAVO
 48    37-58.3N 123-45.1W  NAVO    38-23.2N 122-59.2W  NAVO
DATUM  38-12.2N 123-22.1W          AT  221440Z JUL 75  49.67 HRS DRIFT

COMPUTED TOTAL PROBABLE ERROR OF POSITION (E)  30.6 MILES
STANDARD 3RD SEARCH RADIUS  61.2 MILES
SEARCH AREA 14973 SQ MI
        ....39-13.4N....
124-39.2W         122-05.0W
        ....37-11.0N....

SWEEPWIDTH  2.6 MILES
SUNRISE AT DATUM 1440Z   SUNSET 0115Z

                         DRIFT VECTORS IN KNOTS
          AVERAGE OVER PERIOD                 LAST HOUR
WINCUR            352/0.29                   351/0.28
AVG SEA    147/0.20 MIN 156/0.13 MAX  147/0.20 MIN 198/0.06 MAX
LEEWAY     300/0.94 MIN 000/0.94 MAX  300/0.91 MIN 000/0.91 MAX
MINIMAX TOTAL     337/0.92                   335/0.88
PROBLEM COMPLETE
BT
```

Figure 6-1: SARP Drift Solution.

Although SARP did not relieve the search planner of the plotting burden, it did relieve the search planner of much of the computational burden associated with determining the location and size of the overall search area.

Another program in the computerized SAR system of which SARP was a part provided some additional search planning assistance. This program was called PLAN. The search planner

could input the center, length, width, and orientation of the overall search area and two altitude-sweep width pairs. Then, the search planner could enter up to fifty pairs of SRU speed and on-scene endurance values. Program PLAN would then subdivide the search area into smaller rectangles in such a way that uniform coverage was achieved throughout the search area as a whole. The PLAN output included the center point, length, width, orientation, track spacing and altitude for each sub-area, ensuring that in "Mode 1" (all sub-areas with same orientation) altitudes were alternated to ensure positive altitude separation in adjacent sub-areas.

Yet another program in this system was EPRB, used to estimate the position and the corresponding probable position error of an Emergency Position Indicating Radio Beacon (EPIRB) or Emergency Locator Transmitter (ELT) based on signals acquired by passing aircraft. Inputs included the time of each report, the flight level of the observing aircraft, and the positions where the signal was first acquired, was strongest and was last heard.

Considering the early date, SARP and the several associated ancillary programs were quite sophisticated. SARP's primary limitations were the primitive and occasionally unreliable teletype input/output interface and the coarseness of the environmental data available at the time, except in certain areas like the Florida Straits where the sea current data had an average spatial resolution of about six minutes of latitude and longitude on a monthly basis. A shortcoming of most of the available sea current data was that it was based on climatology on a one-degree spatial scale and either a monthly or seasonal time scale. Exceptions were the Gulf Stream file and the Long Island Sound tidal current files. All of the programs suffered from being implemented on a CDC 3300 mainframe computer that was already obsolete at the time.

Throughout its lifespan of more than 10 years, SARP was maintained and continuously improved by the Operations Analysis Branch of the Information Systems Division of the staff of Commander, Atlantic Area Coast Guard, then located on Governors Island, NY. The program itself proved to be very reliable and it was frequently used until circa 1982 when it was taken out of service with the demise of the CDC 3300 computer on which it was hosted. SARP was not ported to the new PRIME minicomputers located on Governors Island because it was considered obsolete. The Computer Assisted Search Planning (CASP) system, a sophisticated stochastic Monte Carlo-based simulation, had been in operation since about 1974. In 1982, the SAR Program elected to rewrite CASP and implement the new version on the PRIME computers and make it the sole computer-based search planning aid for U.S. Coast Guard search planning. This decision was based in large part on CASP having demonstrated the limitations of manual and automated manual methods.

6.3.2 CANSARP

CANSARP is a direct descendent of the U.S. Coast Guard's original SARP program, at least in concept. However, the current version could not be evaluated directly. Some relatively old CANSARP documentation was reviewed and a few electronic mail messages were exchanged with the Canadian Coast Guard. A demonstration of a new version of CANSARP was seen at the SARSCENE 99 conference in St. John's, Newfoundland, in October 1999. However, answers to many questions about its internal operations were not available. CANSARP runs on Sun Microsystems hardware in a UNIX environment.

CANSARP uses gridded environmental data obtained from Canada's national environmental agency on a near-real-time and forecast basis in a manner similar to the way the USCG CASP system obtains data from the U.S. Navy's Fleet Numerical Meteorology and Oceanography Center in Monterey, California. CANSARP computes eleven datums from a single initial position. Initially, all use the same wind and current data, the only difference being the leeway divergence angle that is chosen. The leeway divergence angles used are the leftmost, rightmost and nine equal divisions in between. Because this number is odd, one trajectory is always directly downwind. As the datums separate, they may encounter differing environmental data over space and time, which means that the eleven datums will form a rough arc, as shown in Figure 6-2 below. The reason the large circle is not tangent to the two "outboard" small circles was one of many unanswered questions. Although not shown in Figure 6-2, CANSARP is supported by a fully capable GIS similar to those of SARMAP, SARIS and C2PC/AMS discussed below.

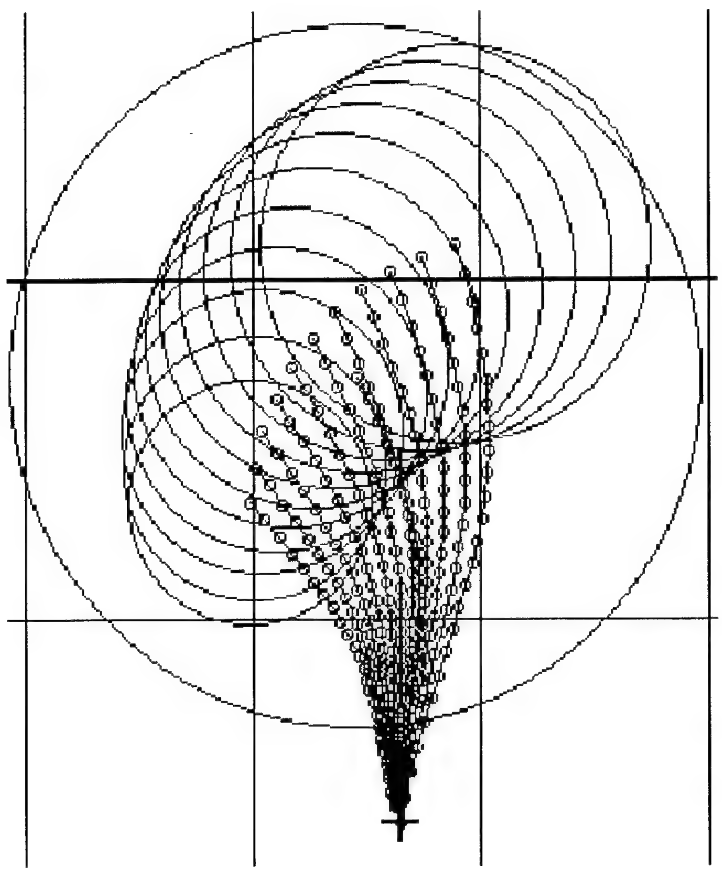

Figure 6-2. CANSARP Drift Solution.

CANSARP re-computes the eleven trajectories from the initial position for all drift intervals. That is, datums used to plan the first day's search are not used as starting points for the second day's drift update. Instead, the entire problem is re-computed from the initial position and time to plan the second day's search. This avoids many of the logical anomalies found in other minimax solutions. CANSARP computes intermediate datums on a one-hour time step. Figure 6-2 shows a 24-hour drift interval with the intermediate datums plotted every hour. This allows CANSARP to make better utilization of gridded data. However, this still does not address the

issue of initial position uncertainties that cover several cells of the environmental data grid. The initial data used for computing drift will come from only the cell containing the estimated initial position, even though there may be a non-trivial probability of the actual initial position being in another cell. Only as a result of the separation of datums over time due to leeway divergence will the number of data cells in use at any one time increase from one to a larger number. Almost all of the data that might affect where the search object drifts, by virtue of different data cells being within the uncertainty region around the datum(s), remains unused.

6.3.3 *ASA* SARMAP/ARCVIEW®

Applied Science Associates, Inc. (ASA) of Narragansett, Rhode Island has developed an automated manual solution that may be packaged with their OILMAP/ARCVIEW® and/or Incident Command System (ICS) software. The purpose of their OILMAP product is to predict oil spill trajectories and perform analyses to assess risks if an oil spill occurs in or near a specific area of interest. Interestingly, this software uses a Monte Carlo approach for estimating the probable distribution of drifting oil over time that is very similar to the way CASP estimates the probable distribution of search object locations over time. However, when *ASA* tried to market this approach for SAR, their foreign clients were not interested because the solution did not appear to match their officially adopted manual methods. As a result, *ASA* was forced to automate a manual method in order to meet client demands. *ASA's* typical method for providing software support for oil spill trajectory modeling is to first obtain or develop detailed environmental data for the client's locale. Since these are in coastal regions, the time and spatial scales need to be quite fine to resolve tides, river outflows, other hydrography, bathymetry and coastal features of interest. These detailed environmental databases are then used as input to their stochastic model of oil spill movement and dispersion. In strategic mode, the model can help planners assess the impacts of various types and sizes of oil spills under a variety of circumstances, including location, date and time, weather conditions, etc. Included in the package is the ability to define an environmental data grid by specifying only a few data points.

SARMAP has many GUI/GIS features that are useful for search planners. As the full name of the software package implies, SARMAP uses the open-architecture ARCVIEW® product as its GIS. Like the Canadian and U.K. tools, SARMAP has access to quite detailed current databases. Wind data are input by the user in increments of as little as one hour without any spatial dependencies. It is easy to input wind data that is constant over larger blocks of time without entering a value for each hour. The software simply fills in the number of one-hour blocks specified by the user from one entry of the wind data. Figure 6-3 shows a SARMAP drift solution.

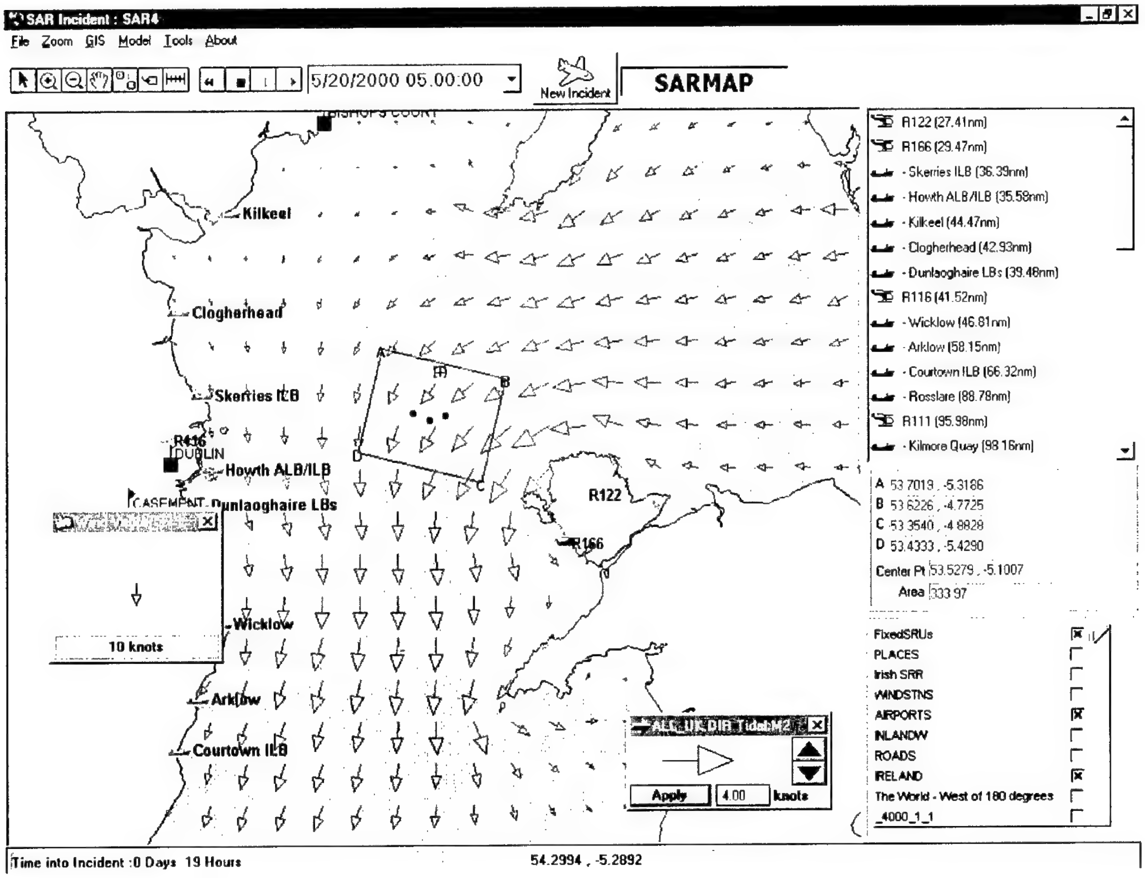

Figure 6-3: SARMAP Drift Solution.

SARMAP computes three datums based on leftmost, rightmost and zero (downwind) leeway divergence, as shown in Figure 6-3 above. The user can vary the time-step between intermediate datums in increments of as little as one minute. The default time-step is one hour but it may be reduced to as little as one minute. When computing drift updates, short time-steps allow it to discern quite small variations in the currents over space and time. Wind data are accessed based on time and is not averaged over the drift interval as in the UK CG3 and USCG C2PC/AMS methods. Unlike the usual minimax methods, SARMAP does not compute a single large minimax circle but circumscribes a single rectangle around the three individual datum search circles. The radii of these circles consist of the total probable error of position computed in the usual manner times the appropriate "safety factor." A rectangle is then circumscribed around the three circles to form the recommended search area.

Of all the tools examined, SARMAP has by far the richest set of search object choices for determining leeway. There are nearly 60 choices presented. These are apparently taken from the leeway taxonomy provided by Allen and Plourde [1999].

SARMAP can also handle both leeway rate and leeway divergence "minimax" solutions simultaneously. Entering two different leeway rate categories or formulas will cause six datums to be produced and the resulting search area will be a rectangle circumscribed around all six search circles. It is assumed that the same leeway divergence applies to both leeway rates, so

this capability is less than completely general. However, SARMAP was the only automated manual solution of those evaluated that could handle this situation at all.

A "stress test" of SARMAP was run using zero currents and reversing winds. Wind data of 270T/30 knots for six hours was input followed by 090T/30 knots for the next six hours. An object with ±35 degrees leeway divergence was used for the test. Animated output showing the computed search object motion was provided, but it moved too fast for viewing on the evaluator's computer. This can no doubt be adjusted in the software. However, a mode was also available to step through the intermediate computed positions one at a time.

During the first six hours, the three datums diverged in an easterly direction as expected and the recommended search area was an increasingly large rectangle with a north-south orientation. During the second six hours, the three datums exactly reversed direction and converged back to the original starting position. This is consistent with the assumption that an object that drifts to the left of the downwind direction will always drift to the left of the downwind direction and similarly for objects that drift to the right of the downwind direction. The search area did not shrink but remained constant in size. However, it changed shape, becoming more and more nearly square until it was exactly square after the datums converged to a single point.

SARMAP takes as much advantage of the detailed gridded environmental data as possible in an automated version of a manual method, primarily by using a short time-step. SARMAP works well, within the basic limitations imposed by the manual technique, and does not seem to suffer from the quirks found in some other systems. *ASA's* implementation uses the datums from one drift interval as starting points for the next interval. The "left" trajectory for leeway divergence always stays to the left of the downwind direction, and similarly for the "right" trajectory. As a result, datums are more likely to converge under conditions of veering, backing, or reversal of wind vectors. *ASA's* method for counteracting the possible decrease in search area size due to converging minimax datums as a result of wind/current shifts is to never let the search area decrease in size from one time-step to the next. That is, the search area size will always be the larger of the previous and currently computed search area sizes. While this is not a "perfect" solution, it is certainly better than producing an unrealistically small search area.

ASA now has detailed current data for most of the U. S. Coast and could be easily adapted to use current data from other sources, such as the offshore products from ocean circulation models. For any maritime search planning tool to work well, especially in coastal areas, accurate and detailed environmental data are essential for both winds and currents.

6.3.4 Search and Rescue Information System (SARIS) with UK CG3 Method

A demonstration of the Search and Rescue Information System (SARIS) developed by BMT Marine Information Systems Limited of Southampton, U.K., was obtained through the kind efforts of HM Maritime and Coastguard Agency Training Centre and BMT. This demonstration consisted of "canned" programs on compact discs designed to show the capabilities of the current operational version (SARIS I) and its soon-to-be-released replacement (SARIS II). There was also a "live" demonstration and a brief opportunity to exercise the SARIS software itself during an informal visit to the Training Centre. In both cases, the search planning methodology implemented in these versions of SARIS was a computerized version of the UK

CG3 manual method. Like *ASA*, *BMT* also produces oil spill trajectory models using Monte Carlo techniques. However, *BMT* was prevented from taking this approach for either SARIS I or SARIS II due to HM Coastguard requirements to use the UK CG3 manual method's technique for computing mean drift trajectories and recommended search areas. Figure 6-4 shows a UK CG3 drift solution as computed and displayed by SARIS.

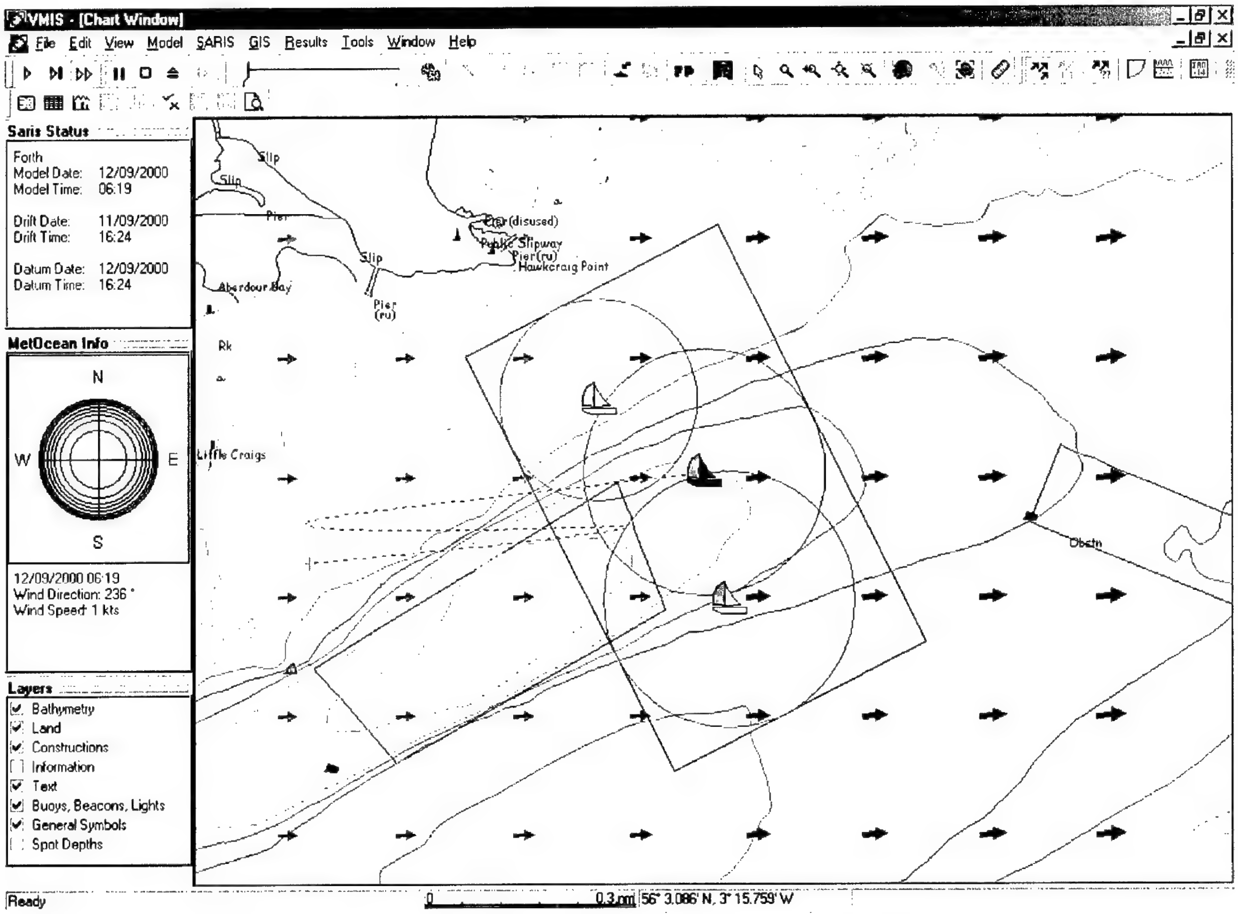

Figure 6-4: UK CG3 Drift Solution in SARIS.

SARIS has many GUI/GIS features that are useful for search planners. In addition to detailed vector shoreline data, SARIS also has detailed vector bathymetry data showing bottom contours. Like the Canadian and *ASA* tool, SARIS has access to quite detailed current databases. These are on a nominal 12-kilometer grid with interpolation algorithms to compute currents for intermediate points. Even finer resolution base data are being considered for areas very near shore and in the smaller bays and estuaries. Wind data are input by the user in 6-hour increments without any spatial dependencies, although an automated source of gridded wind data are being considered as a future enhancement.

The UK CG3 search planning method computes three datums based on leftmost, rightmost and zero (downwind) leeway divergence. SARIS implements this method with a very short 5-minute time-step between intermediate datums when computing drift updates, allowing it to discern quite small variations in the currents over space and time. Unlike some other minimax methods,

the UK CG3 method does not compute a single large minimax circle but circumscribes a single rectangle around the three individual datum "probable error" circles. "Probable error" is in quotes because the U.K. has modified the way these radii are computed. Instead of computing the square root of the sum of the squared probable errors, the U.K. method simply adds the probable errors for the incident position and the drift. This is tantamount to treating probable errors as if they were maximum possible errors. The result is not total probable error in the statistical sense and it should not be referenced by that name.

The U.K. has made some other modifications to the original CSPM methodology as well. No "safety factor" is applied to the radii used to plan the first search rectangle. However, the CSPM "safety factors" are applied for the second and subsequent searches. Another change is that the search craft's probable position error is not used in the usual fashion. Instead of combining it with the other probable errors, the U.K. method treats search craft position error as a maximum, not probable, error. Assigned search areas are enlarged from the desired search areas by a margin large enough to guarantee that the search craft will cover the desired search area if its actual positioning error is less than its maximum positioning error. Usually this results in overlapping search areas when multiple search craft are involved, something most search planners strive to avoid for safety reasons.

In the U.K. documentation, the stated reason for not using a "safety factor" on the first search was that it caused an excessive (21%) increase in the size of the search area. However, the U. K. method effectively treats all probable errors as maximum errors, and adds them in an apparent attempt to obtain a 100 percent probability of containment. This can easily increase the first search area's size by an even greater amount than that obtained by using a correctly computed total probable error multiplied by the first search "safety factor." Furthermore, if the position and drift error values being used really are maximum, and not probable, errors, then "safety factors" should never be used for planning any search. If the probability of containment is already 100 percent, any expansion of the search area will simply waste effort and potentially increase the mean time to find search objects. In practice, UK CG3 solutions are often run for every search in a series as if each were a "first" search, thus effectively removing the use of "safety factors" altogether.

Although the "live" demonstrations were brief and the exercise of the software with inputs specifically chosen to "stress" it even briefer, some problems were observed.

During the "live" demonstration of the UK CG3 method in SARIS II, a POD of 99 percent was computed for a coverage of 1.33. The correct POD for this coverage, according to the POD vs. Coverage curve based on Koopman's inverse cube model of visual detection and parallel track search patterns (the curve that has been in use for many years), a coverage of 1.33 should produce a POD closer to 90 percent.

During the "exercise," the currents were made zero everywhere and two opposing wind vectors of equal magnitudes and durations were input to the UK CG3 method in SARIS I. The initial wind values were 270T/30 knots for six hours followed by 090T/30 knots for another six hours. The search object was chosen to be a life raft without a drogue. According to the present U.K. "operational" SAR manuals, the leeway rate for such objects is given by the formula

$$Leeway = 0.07 \times Wind\ Speed + 0.04$$

and the leeway divergence is ±35 degrees off the downwind direction. The purpose of this test was to isolate the leeway computations, stress the system, and see how the UK CG3 method as implemented in SARIS dealt with the potential problem of converging datums after the wind reversal.

The result was a computed "downwind" leeway in the direction of 270T, or directly ***up***wind, and leeway vectors 35 degrees to the right and left of this direction. A "compass rose" style of display on the right side of the SARIS screen showed the wind as coming from 270T. After six hours of simulated time, this display reversed and showed the wind coming from 090T. However, the computed leeway did not change but continued in a westerly direction that was then appropriate.

The problem was re-run with a wind of 000T/30 knots followed by a wind of 180T/30 knots. This time the computed "downwind" leeway was in the direction of 090T, or at right angles to both of the input wind vectors, with leeway vectors 35 degrees to the right and left of this direction. Again the leeway did not change when the wind data reversed.

These tests indicate that the UK CG3 method implemented in SARIS computes and uses a single average wind vector over the entire drift interval. Since the magnitude of this "average" wind was apparently unaltered in these tests, it appears an incorrect averaging technique is being used, possibly the one that appeared in the U. S. National SAR Manual after a major rewrite in 1973. That edition recommended averaging wind vector data by computing the arithmetic mean of the directions in degrees and using the result as the mean direction. The arithmetic mean of the magnitudes was then computed and used as the mean magnitude. This is an incorrect and ambiguous method for averaging vectors, as the following brief counter-example demonstrates. Assume there are two wind vectors of equal magnitude, one from 350T and the other from 010T. The arithmetic mean of these two values is (350 + 010)/2 = 180T or exactly opposite to the expected mean direction of 000T, assuming the wind veered or backed through 20 degrees and not 340 degrees. This could explain the second result in the preceding paragraph. The first result (upwind leeway) remains inexplicable.

Another question for which there seemed to be no satisfactory answer was why SARIS would average 6-hour wind data over the entire drift interval when the drift computations were being done on a 5-minute time step to take advantage of the detailed current data. It would seem more logical to use the wind vectors "as is" during their respective 6-hour intervals.

Despite not passing the reversing wind "stress test," the UK CG3 method as implemented in SARIS probably computes very good drift trajectories and search areas over the short time intervals typical of most UK SAR cases. As long as wind directions and speeds do not vary too much, and do not pass through 000T, even the suspected incorrect vector averaging technique will not introduce too much error. The main strength of the computerized version of the UK CG3 method clearly lies in the very detailed current data SARIS makes available to it.

Both the Maritime and Coastguard Agency representatives and the *BMT* representative present at the "live" demonstrations agreed that a stochastic Monte Carlo modeling approach would be a far superior tool as compared to the UK CG3 method as implemented in either SARIS version.

6.3.5 USCG C2PC Automated Manual Solution (AMS)

When this software was initially evaluated, a number of quirks and anomalies were discovered. These fell into four categories:

The logical flaws inherent in the minimax technique, especially when applied to more than one drift interval, that were not dealt with effectively,

Lack of documentation about the paradigm on which the software was based, leading to the evaluator entering inappropriate combinations of inputs (which were, nevertheless, accepted by the program and used to produce results that seemed to make no sense),

Inappropriate computational techniques, and

Programming errors.

All problems were reported and documented via a number of electronic mail messages that resulted in corrective action. Not all problems were solved, however. The software will still accept combinations of inputs that are inappropriate to its paradigm and, as a result, will produce nonsense outputs in some cases. Students at the U.S. Coast Guard National SAR School are taught "business rules" for using C2PC/AMS to avoid this shortcoming.

C2PC/AMS is essentially a direct port of the earlier GDOC/AMM (Geographic Display Operations Computer/Automated Manual Method) from the GDOC environment to the C2PC environment. The AMM was developed primarily as a proof-of concept prototype demonstration of the benefits GDOC's GUI/GIS environment could provide search planners. It was not intended to improve upon the manual method in any way in terms of improved drift estimates, improved use of the available data, or improved search plans.

AMS has many GUI/GIS features that are useful for search planners. In addition to detailed vector shoreline data, AMS can also display digitized nautical charts. Unlike the Canadian, *ASA*, and U.K. tools, AMS does not have access to any detailed gridded current databases, although mean tidal current over the drift interval can be obtained from a commercial tidal data product packaged with AMS. The user must enter either a single total water current, or a single tidal current and average sea current for the drift interval and possibly a single "other" current. Currents are assumed to be invariant over space and time. Wind data are input by the user in six-hour increments without any spatial dependencies. A single average surface wind vector is computed from this data for the drift interval. The wind data are used to compute an Ekman local wind current using the Mooney method described in the U. S. *National SAR Manual* [1991]. The mean local wind current for the drift interval is then computed and used. Unlike the other automated versions of the manual method, AMS uses only one time-step – from the beginning of the drift interval to the end – when computing drift updates. It does this by computing and using single mean values for each vector in the problem. Thus AMS presumes a

perfectly homogeneous environment exists everywhere during the drift interval. Therefore, the AMS drift solution will work well only where there are no significant variations in currents or wind within a large radius of the initial position during the drift interval.

The U.S. Cost Guard's Automated Manual Solution (AMS) allows the user to select several variables to which minimax may be applied. Only one minimax variable at a time should be selected. The minimax variable most commonly used is leeway divergence. AMS computes two datums, two probable error circles, one large minimax circle and a circumscribed square search area just as in the manual extension to the CSPM discussed earlier. The environmental data for an initial position is used for the first drift interval. Data for the "min" and "max" datums can be input separately for subsequent drift intervals when the "min" and "max" datums are in separate locations. Such data can also be input for the first drift interval where the "min" and "max" datums are initially at the same location. However, this makes no sense in terms of the software's logic and inappropriate computations result. In fact, it is difficult to make sense of the "min" and "max" labels, especially in the case of leeway divergence and even more especially in drift intervals subsequent to the first. This is because the labels do not actually denote the "min" or "max" of anything. They are simply the names associated with the two datums and their respective drift trajectories. Figure 6-5 shows a C2PC/AMS drift solution.

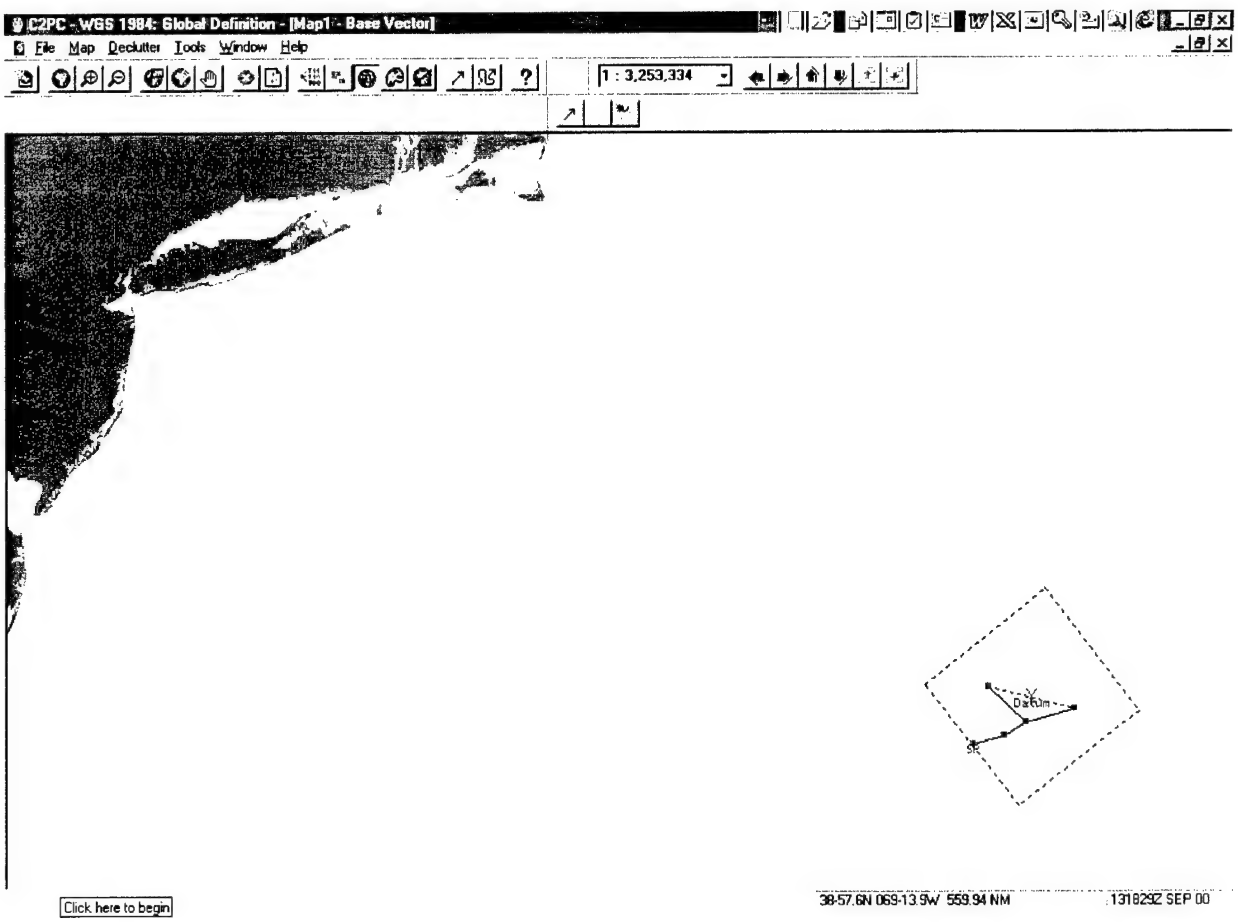

Figure 6-5: C2PC/AMS Drift Solution.

When subjected to the "stress test" of zero current and reversing winds, AMS computed a correct vector average magnitude for the wind of zero. Depending on the orientation of the opposing wind vectors, the associated "direction" was either 000T or 180T. Using the formula for the leeway of a life raft without a drogue, a leeway of 0.04 knots is computed. (*SARMAP* also did this when the entered wind data was zero.) Hence there was a small southerly displacement when the "direction" of the zero wind vector was 000T and a small northerly displacement when the zero wind "direction" was 180T. Unlike *SARMAP*, the search area computed by AMS was very, very small since it computed a drift rate of virtually zero for the entire 12-hour period. This is one of the serious drawbacks of using average values over long distances and/or time periods when drift error is computed as a fraction of the distance drifted.

In the above "stress test," AMS could have been used to compute datums for the end of the first six-hour interval, and then these positions could have been re-entered as initial positions for the second six-hour interval. In this case a quite sizable and more reasonable search area results. When used in this way, AMS addresses the potential problem of decreasing search area size due to the datums turning back toward the initial position as time passes in two ways. First, the probable drift error for the second interval is simply added to the probable drift error of the first interval rather than using the standard formula from statistics. (This technique was already in use in the U.S. manual method prior to AMS.) Second, the software actually computes four datums for the second interval, two for the "min" datum ("Min_min" and "Min_max") from the first interval and two for the "max" datum ("Max_min and "Max_max"). One member of each pair is chosen based on maximizing the distance between the "min" and "max" datums at the end of the second interval. This also means that at the end of the second interval, there is only one "min" datum and one "max" datum. Therefore the same logic may be used again for the third and subsequent intervals. Basically, once the datums have separated in the first drift interval, two minimax solutions are computed in parallel, each with its own environmental data set, and the most widely separated pair at the end of the interval (provided one is chosen from each "column" to prevent both from having the same origin), is then used to determine the next search area and chosen as the "min" and "max" starting points for the next drift interval.

When AMS is used with shorter drift intervals in this fashion, the two simulated search objects may appear to jibe or tack downwind depending on which datums are chosen as "min" and "max" at the end of the interval. Actual jibing behavior has been observed only rarely, if at all, in the leeway experiments done to date. However, the experiments and methods of data analysis may have been biased against revealing instances of jibing. Therefore, the actual frequency of jibing incidents for drifting search objects is presently unknown. Recent simulations conducted at the U.S. Coast Guard R&D Center have revealed that jibing frequency is a critical element in determining the probability density distribution with respect to the downwind direction. A high frequency will lead to a concentration of probability in the downwind direction while even a low frequency of jibing will significantly raise the probability in the downwind direction over what it would be if no jibing occurred. It should be noted that jibing behavior, if it exists, will place search object leeway vectors on the line between the "left" and "right" datums and not beyond that line as simply applying the full leeway rate in the downwind direction would imply and as the CANSARP, *SARMAP* and UK_CG3 tools all presently compute. However, dealing with jibing frequency and many other issues is completely beyond the capabilities of automated versions of manual methods.

Although breaking the desired drift interval into pieces to accommodate a large wind shift produced a more reasonable answer from AMS, it is still disturbing to have the same software give two such vastly different answers to the same problem. This would not happen with *SARMAP*, for example.

The AMS software basically has no "memory" when it comes to tracking the "min" and "max" datums from one interval to the next. The user must copy the "min" datum position and time back into the "min" column of the initial position data sheet and likewise for the "max" datum position. It would be relatively easy for the user to mistakenly reverse these and end up with a mismatch between the datum and the environmental data for that datum. That is, the "min" datum could end up using the "max" datum's environmental data and vice versa. There are a number of other external "business rules" for using AMS that search planners must be taught and must remember in order to use the software correctly. This indicates the software is poorly designed in terms of ergonomics and is therefore prone to incorrect usage.

In coastal regions, a commercial program for computing tidal currents is included as part of the C2PC SAR Tools package. This program uses the standard NOAA models, stations and correction factors. C2PC can import tidal current information from this program and compute the mean net tidal current for the drift interval just as it is done by hand using the worksheets in the *USCG Addendum to the National SAR Manual.*

Of all the automated manual methods examined, C2PC/AMS has to be rated as the poorest among them by a substantial margin. It is the most literal, faithful, translation of the manual method into computer software, and therefore suffers all the limitations of late 1950's pencil-and-paper technology. In terms of the computed drift trajectories, datums and resulting search area recommendations, it is significantly less capable in many important respects than the U.S. Coast Guard's original search and rescue planning (SARP) program developed circa 1970 and abandoned in favor of CASP circa 1982. Although SARP was basically an automated version of the manual method, even then it added some of the more obvious enhancements that a computer makes possible. SARP permitted users to enter wind and current values with times for wind and positions for currents. Intermediate datums were computed every hour and the environmental data values used for the next hour were taken from the data points nearest the intermediate datum time/position. SARP also had access to environmental databases of gridded wind and current data.

In contrast, AMS does not compute intermediate datums, not even for the synoptic wind intervals or tidal cycles. Instead, it computes average values for the initial position over the time adrift and uses these to estimate where the object will go during that time. Variations in the winds and currents over time and space as the drifting object moves along are ignored for up to 48 hours (the maximum AMS drift interval). Early experience with both SARP and CASP showed that the shorter the time step and the higher the spatial and temporal resolution of the environmental data, the better the solution when compared to actual drift trajectories. AMS has no provision for obtaining or using environmental data from automated sources or databases. (Although this capability is planned for oceanic sea current climatology, it is difficult to see how AMS can use such data effectively given its current drift update algorithm.) The user must enter all wind and

non-tidal current values by hand. The AMS implementation of minimax from one drift interval to the next is clearly biased toward making the recommended search area as large as possible. This cannot help but waste large amounts of search effort. When this is combined with the primitive computational method for drift updates, two successive drift intervals of 24 hours each do not necessarily produce the same result as one drift interval of 48 hours. Four successive drift intervals of 12 hours each could produce yet another result in terms of the recommended search area.

The problems just described raise an important issue. Prior to the availability of computer-based search planning tools, the manual method given in the *National SAR Manual* was not a perfectly precise procedure to be followed blindly. Instead, it provided guidance about what data the search planner needed to consider and instructions on specific mathematical procedures such as how to add and average vectors. In some areas, the method was deliberately left somewhat vague or ambiguous to allow the search planner to use his/her experience, judgment, local knowledge, etc. Environmental data was often taken from atlases and charts (e.g., pilot charts) where the search planner could see the approximate structure of currents (and prevailing winds) and how they varied from one place to another in the search object's general vicinity. Thus it was reasonably obvious when a single drift vector based on average values over a long interval was adequate and when it was not. In the latter case, the search planner was allowed, and expected to exercise, the option of using several short drift intervals to account for changing currents and winds as the drift progressed. Such atlases and charts have been largely replaced with digital products that are both less revealing (without an appropriate graphical computer-screen display) and unavailable to C2PC/AMS. This leaves search planners in a difficult position and makes the exercise of their judgment in deciding what data to use and how to use it considerably more difficult. There is the very real possibility that a search planning solution from C2PC/AMS will not be as good as the earlier truly manual solutions were.

In summary, C2PC/AMS does no more than literally automate an imprecise pencil-and-paper methodology. Because software has no judgment capability, the ambiguities of the manual method were replaced with fixed, inflexible interpretations. C2PC/AMS adds little or no value beyond that supplied by the GUI/GIS environment in which it operates. It is just a fast and flashy way to get a solution no better, and quite possibly poorer, than that obtained by truly manual means.

6.3.6 Summary of Automated Manual Solutions

Although all of the automated versions of manual methods provide the search planner with some advantages, none can completely overcome the significant limitations of the pencil-and-paper techniques on which they are ultimately based. Due to their short time steps and improved environmental data, CANSARP, *SARMAP*, and SARIS/UK_CG3 all generally provide increased accuracy in drift trajectory projections over manual methods when the initial position is accurately known, whereas C2PC/AMS does not. Given the constraint of a single starting position and a very limited number of drift trajectories, it appears that CANSARP and *SARMAP* make the best use of the available environmental data and produce the most accurate drift trajectories and best search areas. (SARIS provides very good current data to the UK_CG3 method but that method's use of the supplied wind data are suspect.) In all cases, it is the GUI/GIS environments that support the automated manual solutions, rather than the automated

solutions themselves, that provide the greatest level of support to the search planner over and above the earlier pencil-and-paper methods. It is important to avoid confusing the quality and sophistication of the solution with the quality and sophistication of the packaging it comes in.

6.4 STOCHASTIC COMPUTER SIMULATION

The CSPM and all of its derivatives, including the automated manual solutions, use an analytic approach to the search planning problem. The basic technique in use is one where the mean (average, expected) location of the search object is estimated for the time of the next search, the probable error of that position is also estimated, and the probability density distribution of possible search object locations represented by these two values is assumed to be circular normal. In other words, only two values are computed and used to define the search object location probability density distribution. As we have already discussed, this technique is only valid under the simplest of circumstances. Rarely do the circumstances of actual SAR cases meet the simplicity requirements of this technique.

A much more sophisticated and flexible method is to simulate search object motion (both pre- and post-distress), search facility motion, detection, etc. as stochastic processes. A stochastic process is one that contains some random variation. Stochastic processes generally involve variables whose values are not precisely known or situations where the relationships among variables are not precisely known. For example, search object drift is best modeled as a stochastic process because the values of the environmental forces causing the drift are not precisely known at any given moment and can vary in both predictable and unpredictable ways over space and time, giving a large number of possible combinations. In addition, the object's behavior in response to the environment, e.g., the relationship between leeway and wind, is only approximately understood and cannot be assigned a particular value with absolute certainty. These characteristics of reality lead to a large number of possible drift trajectories. In human terms, simulating stochastic processes requires massive amounts of computations, putting this technique out of reach for manual use. However, today's relatively inexpensive desktop and even laptop computers possess sufficient computing and storage capacity to run quite sophisticated stochastic simulations in reasonable lengths of time, even for the SAR mission with its tight response-time constraints.

The purpose of any simulation or simulator (e.g., a flight simulator) is to artificially mimic reality to the greatest degree of accuracy possible with the available technology and within the time constraints required to make the simulation results useful. Trying to imitate almost any significant aspect of the real world involves a large number of interrelated variables. A simulation uses the known interrelationships among the variables of a problem to build an imitative computer program. In addition, random errors and variations in the values of the variables and the interrelationships among them are also included in the programming since random error and variation are facts of life.

Once programmed, simulations can be used for a variety of purposes. One purpose is predicting the future. Models used to predict the weather are probably the most well known simulations of this type. Simulations can also be used to evaluate different possibilities. In recent years, global climate models have been used to evaluate and predict the effects of different predictions about future levels of "greenhouse gasses" and other variables on global warming. Simulations also

allow ideas and designs to be evaluated in a computer without actually building and testing a complex structure or machine. Simulations are also extensible. Since they are already designed to deal with large numbers of variables and interrelationships, adding another variable and set of relationships to the mix, while not a trivial undertaking, does fit the general paradigm. This is generally not true of simple analytic approximations that deal with few variables. For example, adding jibing frequency to the leeway computations would be a relatively easy enhancement to a stochastic simulation model like CASP, but it would be a very difficult proposition for any analytic model. Attempts to extend the CSPM to deal with additional variables or even just a different distribution of possible values of an existing variable (e.g., leeway rate or leeway divergence) has proven very difficult due to the logical and computational problems with extending the CSPM paradigm to distributions other than circular normal. From the perspective of logical consistency, making such changes to a computer simulation would be trivial in comparison to attempting similar modifications of analytic methods. In short, while simple analytic methods are adequate for simple problems, when solutions or predictions are needed for situations of any significant complexity, the tool of choice is very often computer simulation. Since maritime search planning is often a very complex problem, it is not always amenable to simple solutions. Computer simulation techniques are much more appropriate.

6.4.1 Monte Carlo Modeling (CASP 1.x)

One of the most common and easily explained methods for simulating stochastic processes is the Monte Carlo method. It derives its name from the city of Monaco that is famous for its casino where games of chance are played. A Monte Carlo simulation works by sampling randomly from each individual variable's distribution of possible values and randomly varying the relationships among the variables within known limits. Although it is not an exact analogy, we could say that a variable's exact value within the possible range is based upon a roll of the dice. All samples are independent of one another. The sample values are then combined in an appropriate fashion to generate a sample outcome. If the numbers of samples are large enough, then the distribution of sample outcomes is representative of the real-world distribution of possible outcomes.

Figure 6-6 below shows how a random sample sea current is determined in a Monte Carlo model when the mean sea current and its probable error are known. In Figure 6-6 a mean sea current vector of 030T/1.0 knot is shown with a circular normal distribution of 500 points centered on the "head" of the vector. The probable error of this distribution is 0.3 knots.

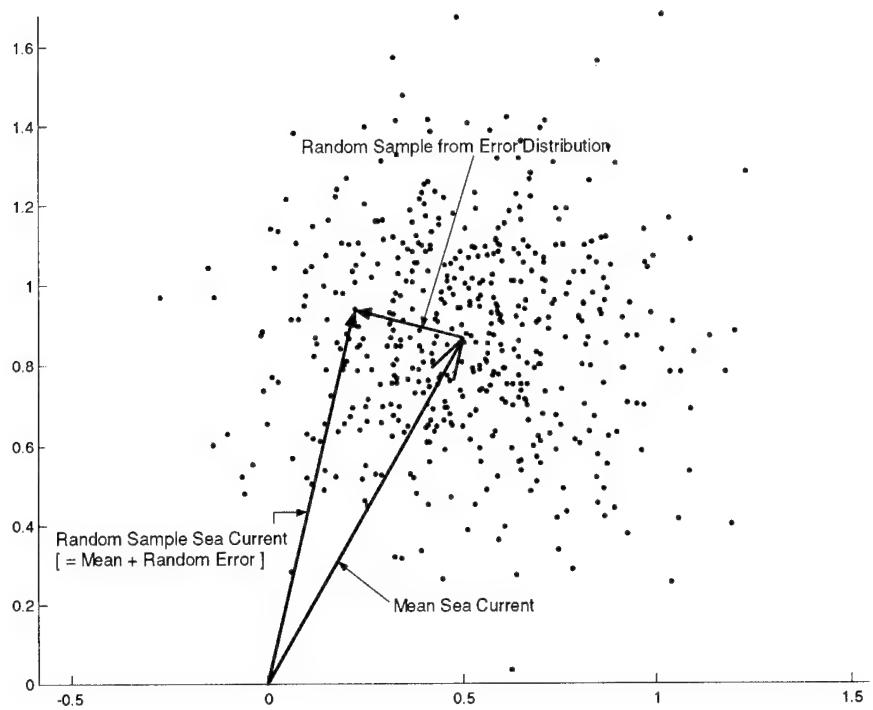

Figure 6-6. Obtaining a Sample Sea Current Vector.

Each point's x- and y-coordinates were chosen independently and at random using a computerized pseudo-random number generator that produces a normal distribution of values. Each dot represents where the "head" of a randomly selected error vector would fall in this particular run of the simulation. We have shown one such error vector. When this error vector is added to the mean sea current vector, the resultant is a sample sea current vector that will be used to compute a sample drift velocity. A sample drift velocity is found by drawing independent samples from the wind, wind current, other currents (if present), and leeway in the same basic fashion and then adding these sample vectors in the usual manner to obtain a sample drift rate and direction. This sample drift velocity is then multiplied by the appropriate amount of time to obtain a sample drift distance from which a sample drift position may be computed. A large number of sample positions computed for a specific point in time will provide a good approximation of the probability density distribution of possible search object locations.

Note: Wind current and leeway have two sources of error. One source is the uncertainty about the exact value of the wind vector used to compute them. The other source is the uncertainty about the relationships between the wind and the leeway and wind current. A sample wind is drawn in exactly the same way as illustrated for sea current in Figure 6-6. Then this sample wind is used to generate wind current and leeway vectors. In each case, an additional error vector is drawn at random from the uncertainties about the relationships between the wind and the leeway and wind current vectors that would still exist even if the wind's exact value were known. The respective error vectors are then added to the leeway and wind current values computed from the sample wind to get sample leeway and wind current vectors.

The U.S. Coast Guard's Computer Assisted Search Planning (CASP) system uses a Monte Carlo technique for simulating search object drift as a stochastic process. CASP initially distributes simulated search objects, called "replications" or "reps" for short, according to parameters entered by the search planner. Three basic distribution types are available and may be used alone or in combination:

Distributions about a point (circular normal, just as in the CSPM),

Distributions along a track line (uniform along the track, normally distributed to either side of track), and

Distributions contained within a specified polygonal area (uniformly distributed throughout the area).

Each replication is "tagged" with the following information:

Present position (latitude, longitude, time of position),

P-fail (probability of non-detection to date),

Last position that was based on analysis ("actual") environmental data vice forecast data and the *P-fail* value at that time,

Status (drifting, temporarily aground, aground, underway) and

Target ID (type of search object it represents), Location ID (the location data set used to specify one of the above types of initial distributions), and Situation ID (identifies the combination of search object and location used to generate the rep).

> *Note: Temporarily aground means a drifting replication encountered land based on a drift trajectory that used forecast environmental data. All drifting and temporarily aground replications have their drift trajectories since the time of the last analysis position re-computed when forecast data are replaced by analysis data.*

Replications are stored sequentially in files in the order that they were initially generated. In outline form, the CASP drift update sequence is:

1. Read the next rep from the input file.
2. Based on the rep's position in space and time, obtain the appropriate environmental data.
3. Draw an independent random sample from each pertinent environmental parameter.
4. Compute wind current and leeway from the sample wind.

5. Draw sample leeway and wind current errors, and compute sample wind current and leeway vectors.

6. Compute the vector sum of the sample leeway and all sample current vectors to obtain a sample drift velocity for this rep.

7. Using this sample drift velocity, compute a new position for this rep for the next whole hour in simulated time, or the desired simulated time for producing a probability map, whichever is earlier.

8. Update the rep's current position tag (latitude, longitude, and time) and, if appropriate, its last analysis position tag.

9. Until the desired simulated time is reached, return to step 1 and repeat this sequence.

Figure 6-7 shows a CASP 1.x drift solution displayed with the aid of the C2PC/CASP interface and GIS. Note that the variation in current speeds from west to east across the Gulf Stream is clearly visible due to the level of detail in CASP's on-line data files. Also note that the distribution does not even remotely resemble a circular normal probability density distribution.

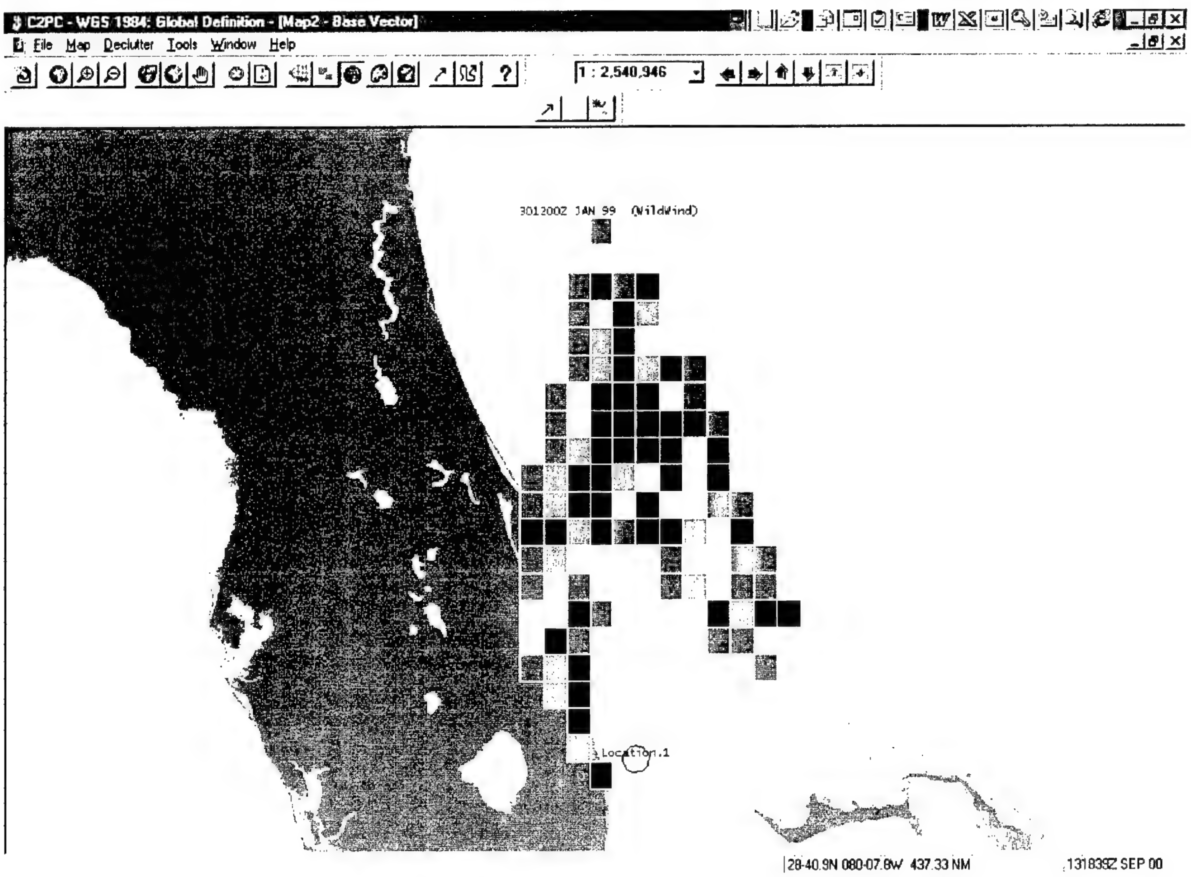

Figure 6-7. CASP 1.x Drift Solution.

The "hotter" colors in Figure 6-7 represent cells with higher probability densities while cooler colors represent cells with lower probability densities. From highest to lowest probability density, the colors are bright red, dark red, orange, yellow, green, blue and gray. The computed probabilities of containment for each of the cells may also be displayed numerically. The ranks of the cells may also be displayed, with cell number one (1) being the one with the highest probability density. Most of the search planning information contained in the CASP 1.x probability map shown in Figure 6-7 would be missing from all of the automated manual solutions. Furthermore, their recommended search areas would almost certainly fail to include some high probability cells while at the same time including some areas with little or even no probability of containing the search object. The result would be wasted search effort and a significantly increased risk of not locating the search object.

When searching has been done, then the distribution is updated for drift as of the mid-search times for each search sub-area, and all the *P-fail* values of reps that are inside the search sub-area at that instant are adjusted according to the POD computed for that search sub-area.

The power of the Monte Carlo method lies in the large numbers of replications it uses. The present version of CASP can accommodate up to 20,000 replications per situation, and a single case can have multiple situations to represent different interpretations of the available data. These situations can also be weighted according to the search planner's assessment of which are more likely to represent the true situation and which are less likely, but still possible. These weights will be reflected on probability maps that display the results of combining the situations to assess where a search object is most likely to be found. The quality of the simulation can always be improved by the simple expedient of adding more replications, until the required computing time becomes excessive.

There are also other advantages. Because of the large numbers of replications, CASP is able to take full advantage of high-resolution gridded environmental data. The CSPM-based methods can, at best, draw environmental data from only one or two (in the case of minimax) locations at one time and in those instances it is restricted to using only the mean value. Automated versions of manual methods can do only slightly better. CANSARP, for example, computes eleven "datums" based on a uniform distribution of leeway divergence angles, so it is possible for CANSARP to use mean values from at most eleven different locations. In all manual and automated manual methods that compute multiple datums from a single starting point, environmental values from different locations will be drawn only after enough time has elapsed for the datums to separate. No matter how much uncertainty there is in an initial position or how much area the initial distribution covers, or how the environmental data may vary from place to place within that area, only mean values for the initial datum can be used initially. CASP, on the other hand, initially distributes tens of thousands of replications in the region of space and time where the incident could have occurred. Each replication draws data from its own individual location in space and time right from the start and continues to do so throughout the simulation.

Another advantage to simulation techniques is flexibility. For example, adding "survivability" to CASP would only require adding another tag to the replications, adding any necessary environmental data (e.g., mean water temperature and its probable error) to the environmental files, and adding the necessary logic to utilize that data (e.g., survivability vs. water temperature

curves). This would also allow CASP to recommend optimal survivor search plans that maximize the probability of finding survivors alive. It is also possible to add the modeling of transitions from one type of search object to another. For example, survivors might stay with a vessel for some time following the distress incident, then abandon ship into a life raft, and then be thrown from the raft into the water by rough seas. This is the multi-state search problem described in Chapter 2. Search craft motion could be added to the CASP model along with more detailed modeling of detection probabilities to show the true effects of the relative motion between search objects and search craft. This would help determine whether a particular subset of the possible drift trajectories was ineffectively covered by a given search. The search planner, once alerted to such a problem, could then plan effective countermeasures if significant amounts of probability were involved. Algorithms that produce optimal search plans for stochastically moving objects (like search objects adrift) are now available and could be added to CASP as an improvement to the current static "snapshot" method of optimization. All of these capabilities are quite beyond those that are possible for any manual or automated manual method.

Taking the flexibility issue a step further, a CASP or CASP-like tool could be modified and improved to support certain non-SAR mission areas of the Coast Guard. For example, some parts of the Coast Guard already have the *ASA OILMAP/ARCVIEW®* as a backup for NOAA software in support of the Marine Environmental Protection (MEP) mission area. The similarities between ASA's (and NOAA's) approach to oil spill trajectory modeling and CASP's modeling of search object trajectories are striking. The very same framework and concepts support both and there is good reason for combining them. The changes needed to make CASP usable in this realm would amount to adding motion models for oil and hazardous chemicals (including appropriate dispersion modeling) and a few MEP-specific inputs. Where "landed" replications are basically discarded for SAR purposes, they would be of primary interest to MEP as they would be used to compute the probabilities of an oil or hazardous chemical spill coming ashore, where, and in what concentrations. SAR would benefit by making CASP usable in the near-shore environment. The Coast Guard as a whole would benefit from a reduced training burden and increased standardization as about 90 percent of the user interface would be the same for both applications. A similar argument can be made for CASP support of law enforcement search and surveillance missions and patrols. (In fact, CASP 1.x does have a "law enforcement" module but it seems to be used only rarely, if at all. Probably its existence is not widely known.) In all cases, a CASP-like tool could be used as both a tactical decision aid to deal with on-going situations and as a strategic aid to plan and evaluate different strategies and evaluate specific tactical responses to various kinds of situations. In particular, an enhanced CASP-like tool could be used to evaluate different search and surveillance patterns, patrol routes, etc.

6.4.2 CASP 2.0

CASP 2.0 could not be evaluated because it was never completed and no functioning version exists. However, Soza & Company, Ltd., performed an evaluation of those CASP 2.0 software modules delivered to the Coast Guard in 1996 [SOZA 1996]. The evaluation showed that CASP 2.0 addressed many of the limitations and shortcomings of CASP 1.x. It also implemented new or extended capabilities. For example, Brown's algorithm for planning optimal searches for moving search objects was implemented. It appeared that the relative motion between search craft and search objects was correctly simulated and that the effects on search effectiveness and post-search probability density distributions were correctly computed.

Unfortunately, CASP 2.0 also had some serious problems of its own. These problems were primarily related to the way the software was developed and the way the project was managed. The greatest difficulty was that the software consisted of modules from an ASW application the developer had already completed. These were modified and cobbled together in ways that made reliable future enhancement and maintenance of the software almost impossible. From a software engineering perspective, CASP 2.0 was a complete "kludge" of poorly matched elements originally intended for other purposes. Nevertheless, the partially completed CASP 2.0 work was a successful proof-of-concept for many needed CASP 1.x improvements. Although CASP 2.0 was never completed, many valuable lessons were learned that will be useful in future development of search planning decision support tools.

6.4.3 Other Simulation Techniques

The Monte Carlo method is not the only technique that may be used in SAR and other applications involving stochastic processes. Another technique uses Markov chains where the transition probabilities for an object in one cell in a grid to move to each of the eight surrounding cells are computed and used to modify probability maps. The Ocean Prediction System, primarily designed to assimilate environmental data from a variety of sources into a single coherent picture, uses objective analysis and an adaptation of the Fokker-Planck equation to predict wind and current statistics from the assimilated data. There are other techniques as well that can be adapted to the search planning and other problems just cited. However, these are generally more esoteric and require more knowledge of higher mathematics to implement and maintain than a Monte Carlo method does. Much of a Monte Carlo simulation is simply repeated applications of a relatively simple model based on a large sampling from the statistics of the various model parameters. For example, anyone familiar with the CSPM would find the CASP 1.x drift update computations for any single drift trajectory quite recognizable.

6.4.4 Matrix Summary of Search Planning Tools

The results of the search planning tool evaluations are summarized in the matrix below. The ratings shown again make it very clear that simulation techniques have many substantial advantages over analytic approaches.

Capability / Tool =>	CASP 1.x	C2PC/AMS	CANSARP	SARMAP	SARIS
Modeling Technique	Monte Carlo	Analytic	Analytic	Analytic	Analytic
Datum Types	Point, Track Line, Area, Combination	Point	Point	Point	Point, Points on a Track Line
Representation of Initial Situation	Tens of thousands of simulated search objects distributed in space and time	Two simulated search objects at one position and time, one with "left" leeway divergence, the other with "right" leeway divergence	Eleven simulated search objects at one position and time with uniformly distributed leeway divergence angles between "left" and "right" values	Three or six simulated search objects at one position and time, with one each for "left," zero, and "right" leeway divergence and up to two leeway rates	Three simulated search objects at one position and time, with one each for "left," zero, and "right" leeway divergence
Drift Updates	Tens of thousands of independent drift trajectories based on sampling from uncertainties, one-hour time-step	Two mean drift trajectories, single time-step equals length of the drift interval	Eleven mean drift trajectories, one-hour time-step	Three or six mean drift trajectories, variable time-step down to one minute	Three mean drift trajectories, five-minute time-step
Use of 2^{nd}-Order Statistics	Excellent	Poor	Poor	Poor	Poor
Use of High-Resolution Environmental Data	Excellent	None	Good	Good	Moderate
Land Recognition	Yes	No	Yes	Yes	Yes
Post-Drift Probability Distributions on Search Object Location	Computed generalized distribution based on search object type(s) and the distributions of environmental data uncertainties.	Assumed circular normal. Locations of centers based on mean environmental data values. Probable error based on distance drifted.	Assumed circular normal. Locations of centers based on mean environmental data values. Probable errors based on distance drifted.	Assumed circular normal. Locations of centers based on mean environmental data values. Probable error based on distance drifted.	Assumed circular normal. Locations of centers based on mean environmental data values. Probable error based on distance drifted.

Capability / Tool =>	CASP 1.x	C2PC	CANSARP	ASA	SARIS
Probability Maps	Yes	No	No	No	No
Optimal Search Plans	Yes (Static)	No	No	No	No
Computed POS of Search Plan (Search Effectiveness)	Yes	No	No	No	No
Computed Cumulative POS (Effectiveness of all Searching to Date)	Yes	No	No	No	No
Proper Use of Previous Search Results to Plan Future Searches	Yes	No	No	No	No
Accounts for Effects of Relative Motion on Search Effectiveness	No	No	No	No	No
Multiple Weighted Scenarios/Situations	Yes	No	No	No	No
Adaptability to Non-SAR USCG Missions	Optimal tactical and strategic search and surveillance plans for law enforcement, hazardous chemical spill trajectory prediction, and risk analysis.	No significant potential benefit to other USCG missions	No significant potential benefit to other USCG missions	No significant potential benefit to other USCG missions	No significant potential benefit to other USCG missions

6.5 OTHER SEARCH PLANNING DECISION SUPPORT TOOLS

Other tools mentioned in the statement of work were the Ocean Prediction System (OPS) and the Search Master/Case Master system under development in Canada. These tools were not evaluated because neither currently has a search planning module.

The primary purpose of OPS is to assimilate environmental data from a variety of sources to create a meaningful, accurate, and comprehensive environmental picture for search planning and other tools to use. While some of its techniques could be adapted to search planning, a detailed analysis of how this might best be done was beyond the scope of this project.

The Search Master/Case Master tool was viewed at Canada's SARSCENE 99 conference in St. John's, Newfoundland at the developer's vendor booth during the conference. There it was learned that no search planning modules were yet in place. However, it was planned to interface CANSARP with Search Master/Case Master as the search planning module. The possibility of similar interfaces to other search planning tools, based on client desires, was left open.

6.6 CASE MANAGEMENT TOOLS

No case management tools were provided for evaluation, but brief demonstrations of two such tools were obtained. Both were adaptations of COTS software packages. One of the tools, Search Master/Case Master, had already been evaluated by the U.S. Coast Guard Atlantic Area Command Center.

There is no doubt that Command Centers/RCCs and the U.S. Coast Guard as a whole could benefit greatly from having a standard electronic case management system in place. Maximum benefit would be obtained by using it for all types of "cases," not just SAR. The Atlantic Area evaluation of Case Master was very positive. Their concurrent evaluation of the C2PC/AMS software was clearly negative, mainly because of its poorly designed user interface and its complete lack of case management capabilities. Because Case Master is already set up to aid in SAR case management, reporting and data collection, the modifications needed to adapt it for USCG SAR and other mission area case management needs are minimal.

The other case management tool investigated is under development by DoD's Joint Personnel Recovery Agency. It is based on the GEM® product from Electronic Information Systems, Inc. This COTS software is being tailored to meet specific DoD needs in connection with recovering personnel under hostile conditions, i.e., combat SAR. It is being designed to work with DoD databases, some of which may be classified, that contain information pertinent to recovery operations. Specifically, it is touted as an information management system for personnel recovery. It helps the mission coordinator collect, organize, filter and distribute information. At present, it is geared primarily toward recovery of aircrews shot down behind enemy lines, particularly fighter pilots. Types of information include identifying features of the downed pilot, the pilot's evasion plan of action (a completed form left on file), management of isolated personnel reports (ISOPREPs), authentication data (e.g., question, answer sequences to ensure the person recovered is not an enemy in the missing person's clothes), etc. Although the basic GEM software could be adapted to meet U.S. Coast Guard needs, the DoD personnel recovery management version is not suitable for U.S. Coast Guard use.

CHAPTER 7.

CONCLUSIONS AND RECOMMENDATIONS

This report describes the developments in the field of search theory from its origins during World War II to the present day. The principles of search theory were first applied to SAR planning when an unclassified version was published around 1957. To apply search theory to practical SAR planning problems, since computers were not then widely available, compromises and simplifications to search theory had to be made to develop a method feasible for hand calculations. This became known as the "classical search planning method" (CSPM), and it remains the basis for search planning support tools today. The Search and Rescue Planning (SARP) system, the first implementation of search planning support on computers, occurred around 1970, well before the microcomputer age. SARP was basically a computerized version of the then current version of the modified CSPM with somewhat improved use of environmental data and drift computations over purely manual techniques. A few years later SARP was joined by the Computer Assisted Search Planning (CASP) system that took a computer simulation approach to the search planning and evaluation problem.

Unfortunately, Coast Guard SAR search planning support tools have not kept up with technological advances in three important respects. First, they have not kept up with advances made in search theory and algorithm development that are relevant to the practical application of search theory using computer simulation. Substantial advances in search theory applicable to SAR and other U.S. Coast Guard missions have been made in the 25 years since CASP's initial implementation, but little has been put into practice. Advances include Brown's algorithm for moving search objects and Stone's Generalized Search Optimization (GSO) technique that can also accommodate changes in the state of the search object. The oversimplified nature of manual and automated manual methods does not permit any significant advances or benefits to be realized in this area.

Second, modifications made to the CSPM over time to make it applicable to typical, complex search scenarios are inconsistent with the basic theory. The original CSPM was a scientifically based, carefully designed, and well-integrated analytic method that was appropriate for the technology and data available at the time of its implementation in 1957. However, due to the technological limitations of the time, it could adequately address only the simplest of SAR incident scenarios. A few restrictions to keep the problem simple include a known incident position and time in an area of constant wind and current with a specific level of search effort applied in a perfect pattern relative to the drifting survivors. Later, attempts were made to extend the methodology to more typical, and more complex, situations. Unfortunately, the connection between the basic methodology and the principles of search theory seems to have been lost sometime after 1963. As a result, a number of sub-optimal "field modifications" were made that are inconsistent with the underlying theory. The most notorious of these modifications is the Min/Max technique, which violates some of the basic underlying assumptions of the CSPM and the scientific principles on which it was based. Partial attempts to rectify this situation, like the mid-point compromise, were not always improvements.

Third, and most importantly, the tools have generally not kept pace with the significant increases in the amount, level of detail, or accuracy of environmental data products or with new knowledge about drift behavior (particularly leeway) or detection. Manual methods and their automated versions are poorly matched to the large amounts, high quality and detailed resolution of environmental, detection, leeway behavior and other data now readily available. While CASP is reasonably up-to-date with respect to offshore environmental data products, it has never had access to the data it needs in the coastal environment where most SAR cases occur. Likewise, CASP does not contain the latest leeway parameters nor does it make best use of the detection information obtained in the Coast Guard's sweep width experiments.

These shortcomings are evident in the four current computerized SAR planning support tools (U.S., Canadian, British) based on the manual method. All suffer from the limitations of the CSPM and its subsequent modifications in terms of the search plans they can provide. The Coast Guard's C2PC/SAR Tools Automated Manual Solution suffers the most in this regard because it is the most faithful to the manual technique. The other three similar tools provide considerable improvement in the availability and use of environmental data for drift computations. Despite this, none of these manual-based tools can make reasonably complete use of today's detailed environmental data products when compared to computer simulation techniques such as the Monte Carlo method employed by CASP.

CASP 1.x is the most capable search planning tool available in the world today, but only because it is the only one that even attempts to use simulation techniques. The basic framework of CASP's design is more than 25 years old and even then, it was severely limited by the obsolete computing environment in which CASP was forced to operate. As a result, many of the benefits of simulation technology have not been realized. In short, CASP 1.x is a simulation, but it is a primitive implementation of a sophisticated methodology and many key elements are still missing.

In more recent years, the greatly increased availability of inexpensive but powerful microcomputers with high-resolution color geographic information systems has made sophisticated near-real-time computing support both a reality and a normal expectation. Merely, programming powerful workstations to emulate the hand calculations of the manual method—a technique developed to get around the lack of computing capability in the 1950's—is unconscionable. Nevertheless, the Coast Guard currently uses more sophisticated techniques to predict oil spill trajectories and perform risk analyses than for matters of life and death.

The chances of a survivor's continued survival decrease rapidly with time while the risk to searchers and the cost of the search increase with time. Therefore, the primary objective of any search plan is to maximize the chances for finding the survivors sooner rather than later with the available search resources. An optimal search plan is one that produces the maximum probability of success in the minimum amount of time. Based partly on CASP's success, the U. S. Navy eventually developed a program to support search and surveillance operations to monitor Soviet submarine activity that incorporated later developments in search theory and related algorithms. Analysis of the results obtained from using this program showed a remarkable doubling of search effectiveness as compared with previous methods using exactly the same search platforms, sensors, and crews. That is, using a search theory approach to their

tactical deployment and use made the same resources twice as effective. This kind of result shows the importance of maintaining a scientifically sound approach to search planning. By continuing to use a search method that contains the above mentioned errors and inaccuracies, the Coast Guard has decreased effectiveness while increasing costs.

It is recommended that the Coast Guard's SAR planning theory be corrected and that stochastic analysis be the primary method for executing search plans. The U. S. Coast Guard needs and deserves a new computer simulation-based search planning support tool that takes full advantage of the advances in these areas to ensure efficient, effective use of expensive search resources. Those awaiting rescue deserve the time advantage such a planning support tool can offer.

REFERENCES

Allen, A. A. and Plourde, J. V. (1999) "Review of Leeway: Field Experiments and Implementation." U.S. Coast Guard Research and Development Center Report No. CG-D-08-99, April 1999.

ASWORG (1942) "Preliminary Report on the Submarine Search Problem" by the ASWORG 1 May.

Auger, J. M. (1991) An Infiltration Game on *k* Arcs. *Naval Research Logistics* **38**, 511-529.

Benkoski, S.J. (1978) *CAS Operational Summary*, COMPATWINGSPAC Memorandum 15 May (Secret Noforn).

Benkoski, S. J., Monticino, M. G., and Weisinger, J. R. (1991) A Survey of the Search Theory Literature. *Naval Research Logistics* **38**, 469–494.

Brown, S. S. (1980) Optimal Search for a Moving Target in Discrete Time and Space. *Operations Research* **28**, 1275–1289.

Charnes, A. and Cooper, W. W. (1958) The Theory of Search: Optimum Distribution of Search Effort. *Management Science* **5**, 44–50.

DeGuenin, J. (1961) Optimum Distribution of Effort: An Extension of the Koopman Basic Theory. *Operations Research* **9**, 1–7.

Discenza, J. H. (1980) A Solution for the Optimal Multiple Rectangle Problem, in *Search Theory and Applications*, K. B. Haley and L. D. Stone (eds.). Plenum Press, New York.

Discenza, J. H. and Stone, L. D. (1981) Optimal Survivor Search with Multiple States *Operations Research* **29**, 309–323.

Eagle, J. N. (1984) Optimal Search For a Moving Target when the Search Path is Constrained. *Operations Research* **32**, 1107–1115.

Eagle, J. N. and Yee, J. R. (1990) An Optimal Branch and Bound Procedure for the Constrained Path, Moving Target Search Problem. *Operations Research* **38**, 110–114.

Eagle J. N. and Washburn, A. R. (1991) Cumulative Search-Evasion Games *Naval Research Logistics* **38**, 495–510.

Gal, S. (1980) *Search Games*, Academic Press.

Hufford, G. L., and Broida, S. (1974). "Determination of Small Craft Leeway." U.S. Coast Guard Research and Development Center Report No. 39/74, December 1974.

International Aeronautical and Maritime Search and Rescue (IAMSAR) Manual (1999), International Maritime Organization/International Civil Aviation Organization, London/Montreal, 1999.

James, R. (1966). *Ocean Thermal Structure Forecasting*, SP-105, ASWEPS MANUAL, Vol. 5, Washington, D. C., 1966.

Koopman, B. O. (1946) *Search and Screening*. Operations Evaluations Group Report No. 56, Center for Naval Analyses, Alexandria, VA.

Koopman, B. O. (1980) *Search and Screening: General Principles with Historical Applications*. Pergamon Press, New York.

Koopman, B. O. (1956a) The Theory of Search, Part I: Kinematic Bases. *Operations Research* **4**, 324–346.

Koopman, B. O. (1956b) The Theory of Search, Part II: Target Detection. *Operations Research* **4**, 503–531.

Koopman, B. O. (1957) The Theory of Search, Part III: The Optimum Distribution of Searching Effort. *Operations Research* **5**, 613–626.

McCoy, K. F. (1978) Operational *ASW Search Information System (OASIS)*, COMPATWINGSPAC Memorandum 7 March (Confidential).

Morse, P. M. (1982) In Memoriam: Bernard Osgood Koopman, 1900–1981. *Operations Research* **30**, 417–427.

National Search and Rescue Manual (1991), Interagency Committee for Search and Rescue, Washington, DC, 1991.

Newton, F. (1986) *Operations Plan: Final Report,* Triton Technology Report to Recovery Ltd. dated May 22.

Richardson, H. R. (1967) "Operations Analysis" Chapter V of Part 2 in *Aircraft Salvage Operation, Mediterranean*; Report to Chief of Naval Operations from the Supervisor of Salvage and the Deep Submergence Systems Project, published by the US Navy.

Richardson, H. R. and Discenza, J. H. (1980) The United States Coast Guard Computer-Assisted Search Planning System (CASP). *Naval Research Logistics* **27**, 659–680.

Richardson, H. R. and Stone, L. D. (1971) Operations Analysis during the Underwater Search for Scorpion. *Naval Research Logistics* **18**, 141–157.

Soza & Co., Ltd. (1996), Search Planning Methods Analysis, Unpublished report for U. S. Coast Guard Office of Search and Rescue (G-OPR), 200 pages.

Sternhell and Thorndike (1946) *Antisubmarine Warfare in World War II*. Operations Evaluations Group Report No. 51, Center for Naval Analyses, Alexandria, VA.

Stewart, T. J. (1979) Search for a Moving Target when Searcher Motion is Restricted. *Computers and Operations Research* **6**, 129–140.

Stewart, T. J. (1980) Experience with a Branch and Bound Algorithm for Constrained Searcher Motion, in *Search Theory and Applications*, K. B. Haley and L. D. Stone (eds.). Plenum Press, New York.

Stewart, T. J. (1981) A Two-Cell Model of Search for an Evading Target, *European Journal of Operations Research*, **8** 369–378.

Stone, L. D. (1979) Necessary and Sufficient Conditions for Optimal Search Plans for Moving Targets. *Mathematics of Operations Research* **4**, 431–440.

Stone, L. D. (1989) *Theory of Optimal Search*, second edition. Published by INFORMS, Linthicom, MD.

Stone, L. D. (1992) Search for the *SS Central America*: Mathematical Treasure Hunting. *Interfaces* **22**, 32–54.

Stromquist, W. R. and Stone, L. D. (1981) Constrained Optimization of Functionals with Search Theory Applications Mathematics of Operations Research **6**, 518–529.

Washburn, A. R, (1980) Search-Evasion Game in a Fixed Region, *Operations Research*, **28** 1290–1298.

Washburn, A. R, (1983) Search for a Moving Target: The FAB Algorithm, *Operations Research*, **31** 739–751.